
Flow-based reputation with uncertainty

Evidence-Based Subjective Logic

Boris Škorić, Sebastiaan J.A. de Hoogh, and Nicola Zannone

Abstract The concept of reputation is widely used as a measure of trustworthiness based on ratings from members in a community. The adoption of reputation systems, however, relies on their ability to capture the actual trustworthiness of a target. Several reputation models for aggregating trust information have been proposed in the literature. The choice of model has an impact on the reliability of the aggregated trust information as well as on the procedure used to compute reputations. Two prominent models are flow-based reputation (e.g., EigenTrust, PageRank) and Subjective Logic based reputation. Flow-based models provide an automated method to aggregate trust information, but they are not able to express the level of uncertainty in the information. In contrast, Subjective Logic extends probabilistic models with an explicit notion of uncertainty, but the calculation of reputation depends on the structure of the trust network and often requires information to be discarded. These are severe drawbacks.

In this work, we observe that the ‘opinion discounting’ operation in Subjective Logic has a number of basic problems. We resolve these problems by providing a new discounting operator that describes the flow of evidence from one party to another. The adoption of our discounting rule results in a consistent Subjective Logic algebra that is entirely based on the handling of evidence. We show that the new algebra enables the construction of an automated reputation assessment procedure for arbitrary trust networks, where the calculation no longer depends on the structure of the network, and does not need to throw away any information. Thus, we obtain the best of both worlds: flow-based reputation and consistent handling of uncertainties.

Keywords Reputation systems · Evidence Theory · Subjective Logic · Flow-based reputation models

1 Introduction

Advances in ICT and the increasing use of the Internet have resulted in changes in the way people do everyday things and interact with each other. Everything people do happens online, significantly increasing the number of business transactions carried out daily over the Internet. Often, users have to decide whether to interact with services or users with whom they have never interacted before. Uncertainty about services and users’ behavior is often perceived as a risk [2] and, thus, it can restrain a user from engaging in a transaction with unknown parties. Therefore, to fully exploit the potential of online services, platforms and ultimately online communities, it is necessary to establish and manage trust amongst the parties involved in a transaction [11, 40, 43].

Reputation is widely adopted to build trust among users in online communities where users do not know each other beforehand. The basic idea underlying reputation is that a user's past experience as well as the experience of other users influences his decision whether to repeat this interaction in the future. Thus, reputation provides an indication of services' and users' trustworthiness based on their past behavior [36]. When a user has to decide whether to interact with another party, he can consider its reputation and start the transaction only if it is trustworthy. Therefore, a reputation system, which helps managing reputations (e.g., by collecting, distributing and aggregating feedback about services and users' behavior), becomes a fundamental component of the trust and security architecture of any online service or platform [42].

The application and adoption of reputation systems, however, relies on their ability to capture the actual trustworthiness of the parties involved in a transaction [41]. The quality of a reputation value depends on the amount of information used for its computation [10, 17, 27]. A reputation system should use "sufficient" information. However, it is difficult to establish the minimum amount of information required to compute reputation; also, different users may have a different perception based on their risk attitude [2]. For instance, some users may accept to interact with a party which has a high reputation based on very few past transactions, while other users might require more evidence of good behavior. Therefore, a reputation system should provide a level of confidence in the computed reputation, for instance based on the amount of information used in the computation [15, 32]. This additional information will provide deeper insights to users, helping them decide whether to engage in a transaction or not. In addition, reputation systems should provide an effective and preferably automated method to aggregate the available trust information and compute reputations from it.

Reputation systems usually rely on a mathematical model to aggregate trust information and compute reputation [13]. Several mathematical models for reputation have been proposed in the literature. These models can be classified with respect to the used mathematical foundations, e.g. summation and averaging [14], probabilistic models [15, 21, 28], flow-based [6, 23, 25, 35], fuzzy metrics [3, 38]. As pointed out in [42], the choice of the type of model has an impact on the type and amount of trust information as well as on the procedure used to compute reputation.

Among the others, two prominent reputation models are the flow-based model and Subjective Logic (SL) [15]. Flow-based reputation models use Markov chains as the mathematical foundation. Flow-based models provide an automated method to aggregate all available trust information. However, they are not able to express the level of confidence in the obtained reputation values. On the other hand, SL is rooted in the well-known Dempster-Shafer theory [34]. SL provides a mathematical foundation to deal with opinions and has the natural ability to express uncertainty explicitly. Intuitively, uncertainty incorporates a margin of error into reputation calculation due to the (limited) amount of available trust information. SL uses a *consensus* operator ' \oplus ' to fuse independent opinions and a *discounting* operator ' \otimes ' to compute trust transitivity. This makes SL a suitable mathematical framework for handling trust relations and reputation, especially when limited evidence is available. However, the consensus operator is rooted in the theory of evidence, while the discounting operator is based on a probabilistic interpretation of opinions. The different nature of these operators leads to a lack of "cooperation" between them. As a consequence, the calculation of reputation depends on the shape of the *trust network*, the graph of interactions, in which nodes represent the entities in the system and edges are labeled with opinions. Depending on the structure of the trust network, some trust information may have to be discarded to enable SL-based computations.

Our desideratum is to have a reputation system which has the advantages of both flow-based reputation models and SL. In particular, the goal of this work is to devise the mathematical foundation for a flow-based reputation model with uncertainty. We make the following contributions towards this goal:

- We observe that the discounting rule ‘ \otimes ’ in SL does not have a natural interpretation in terms of evidence handling. We give examples of counterintuitive behavior of the \otimes operation.
- We present a brief inventory of the problems that occur when one tries to combine SL with flow-based reputation metrics.
- We present a simplified justification of the mapping between evidence and opinions in SL.
- We introduce a new scalar multiplication operation in SL which corresponds to a multiplication of evidence. Our scalar multiplication is consistent with the consensus operation (which amounts to addition of evidence), and hence satisfies a distribution law, namely $\alpha \cdot (x \oplus y) = (\alpha \cdot x) \oplus (\alpha \cdot y)$.
- We introduce a new discounting rule \boxtimes . It represents the *flow of evidence* from one party to another. During this flow, lack of trust in the party from whom evidence is received is translated into a reduction of the amount of evidence; this reduction is implemented using our new scalar multiplication rule. Our new discounting rule satisfies $x \boxtimes (y \oplus z) = (x \boxtimes y) \oplus (x \boxtimes z)$. This right-distribution property resolves one of the problems of SL.

Our new discounting rule multiplies evidence instead of opinions and is thus fully based on the handling of evidence. In contrast to the old discounting rule, the new one does not have associativity. This, however, does not pose a problem since the flow of evidence has a well defined direction. The adoption of our discounting rule results in an opinion algebra that is entirely centered on the handling of evidence. We refer to it as **Evidence-Based Subjective Logic** (EBSL).

We show that EBSL provides a solid foundation for the development of reputation models able to express the level of confidence in computed reputations. Moreover, having an opinion algebra rooted in a single foundation makes it possible to define an automated procedure to compute reputation for arbitrary trust networks. We believe that having an automated procedure is a critical and necessary condition for the implementation and adoption of reputation systems in online services. To demonstrate the applicability of EBSL to the development of reputation systems, we make the following contributions:

- We show that replacing SL’s discounting operation \otimes by the new \boxtimes solves all the problems that usually occur when one tries to combine flow-based reputation with SL, in particular the problem of evidence double-counting and the ensuing necessity to make computations graph-dependent and to discard information.
- Using EBSL, we construct a simple iterative algorithm that computes reputation *for arbitrary trust networks* without discarding any information. Thus, we achieve our desideratum of automated computation of flow-based reputation with uncertainty.

We stress that this is only one out of many possible reputation models that can be constructed on top of EBSL. EBSL can be used as a foundation for other existing reputation models, e.g. reputation models based on random walks [9]. The investigation of these models is an interesting direction for future work.

The remainder of the paper is organized as follows. The next section presents an overview of flow-based reputation models and SL. Section 3 discusses the limitations of SL, and Section 4 illustrates these limitations when combining flow-based reputation and SL. Section 5

revisits SL and introduces the new EBSL scalar multiplication and discounting operators. Section 6 presents our flow-based reputation model with uncertainty along with an iterative algorithm that computes reputation for arbitrary trust networks. Section 7 presents an evaluation of our approach using both synthetic and real data. Finally, Section 8 discusses related work, and Section 9 concludes the paper providing directions for future work.

2 Preliminaries

In this section we present an overview of flow-based reputation models (based on [35]) and of Subjective Logic (based on [15]). We also introduce the notation used in the remainder of the paper.

2.1 Flow-based reputation

Flow-based reputation systems [6, 23, 25, 35] are based on the notion of transitive trust. Intuitively, if an entity i trusts an entity j , it would also have some trust in the entities trusted by j .

Typically, each time user i has a transaction with another user j , she may rate the transaction as positive, neutral, or negative. In a flow-based reputation model, these ratings are aggregated in order to obtain a Markov chain. The reputation vector (i.e., the vector containing all reputation values) is computed as the steady state vector of the Markov chain; one starts with a vector of initial reputation values and then repeatedly applies the Markov step until a stable state has been reached. This corresponds to taking more and more indirect evidence into account.

Below, we present the metric proposed in [35] (with slightly modified notation) as an example of a flow-based reputation system.

Example 1 Let A be a matrix containing aggregated ratings for n users. It has zero diagonal (i.e., $A_{ii} = 0$). The matrix element A_{ij} (with $i \neq j$) represents the aggregated ratings given to j by i . Let $\mathbf{s} \in [0, 1]^n$, with $\mathbf{s} \neq \mathbf{0}$, be a ‘starting vector’ containing starting values assigned to all users. Let $\alpha \in [0, 1]$ be a weight parameter for the importance of indirect vs. direct evidence. The reputation vector $\mathbf{r} \in [0, 1]^n$ is defined as a function of α , \mathbf{s} and A by the following equation:

$$r_x = (1 - \alpha)s_x + \alpha \sum_{y=1}^n \frac{r_y}{\ell} A_{yx} \quad (1)$$

where $\ell = \sum_{z=1}^n r_z$.

Eq. (1) can be read as follows. To determine the reputation of user x we first take into account the direct information about x . From this we can compute s_x , the reputation initially assigned to x if no further information is available. However, additional information may be available, namely the aggregated ratings in A . The weight of direct versus indirect information is accounted for by the parameter α . If no direct information about x is available, the reputation of x can be computed as $r_x = \sum_y (r_y / \ell) A_{yx}$, i.e. a weighted average of the reputation values A_{yx} with weights equal to the normalized reputations. Adding the two contributions, with weights α and $1 - \alpha$, yields (1): a weighted average over all available information.

The equation for solving the unknown \mathbf{r} contains \mathbf{r} . A solution is obtained by repeatedly substituting (1) into itself until a stable state has been reached. This is the steady state of the Markov chain. It was shown in [35] that Eq. (1) always has a solution and that the solution is unique.

Intuitively, A can be seen as an adjacency matrix of a trust network where nodes represent entities and edges represent the direct trust that entities have in other entities based on direct experience. Based on the results presented in [35], Eq. (1) can be applied to assess reputation for arbitrary trust networks.

2.2 Subjective Logic

Subjective Logic (SL) is a trust algebra based on Bayesian theory and Boolean logic that explicitly models uncertainty and belief ownership. In the remainder of this section, we provide an overview of SL based on [15].

The central concept in SL is the three-component *opinion*.

Definition 1 (Opinion and opinion space) [15] An opinion x about some proposition P is a tuple $x = (x_b, x_d, x_u)$, where x_b represents the belief that P is provable (*belief*), x_d the belief that P is disprovable (*disbelief*), and x_u the belief that P is neither provable nor disprovable (*uncertainty*). The components of x satisfy $x_b + x_d + x_u = 1$. The space of opinions is denoted as Ω and is defined as $\Omega = \{(b, d, u) \in [0, 1]^3 \mid b + d + u = 1\}$.

An opinion x with $x_b + x_d < 1$ can be seen as an incomplete probability distribution. In order to enable the computation of expectation values, SL extends the three-component opinion with a fourth parameter ‘ a ’ called relative atomicity, with $a \in [0, 1]$. The probability expectation is $E(x) = x_b + x_u a$. In this paper, we will omit the relative atomicity from our notation, because in our context (trust networks) it is not modified by any of the computations on opinions. In more complicated situations, however, the relative atomicity is modified in nontrivial ways.

Opinions are based on evidence. Evidence can be represented as a pair of nonnegative finite numbers (p, n) , where p is the amount of evidence supporting the proposition, and n the amount that contradicts the proposition. The notation $e = p + n$ is used to denote the total amount of evidence about the proposition. There is a one-to-one mapping between an opinion $x \in \Omega$ and the evidence (p, n) on which it is based,

$$(x_b, x_d, x_u) = \frac{(p, n, 2)}{p + n + 2}; \quad (p, n) = \frac{2(x_b, x_d)}{x_u}. \quad (2)$$

The bijection (2) holds for any value of the atomicity a . It has its origin in an analysis of the a posteriori probability distribution (Beta function distribution) of the biases which underlie the generation of evidence [15]. This distribution is given by

$$\rho(t|p, n, a) = \frac{t^{-1+p+2a}(1-t)^{-1+n+2(1-a)}}{B(p+2a, n+2-2a)}, \quad (3)$$

where t is the probability that proposition P is true, and $B(\cdot, \cdot)$ is the Beta function. The opinion $x = \frac{(p, n, 2)}{p + n + 2}$ is based on the vague knowledge (3) about t . The left part of (2) holds because the thus constructed opinion x has the same expectation as (3).

Intuitively, the amount of positive and negative evidence about a proposition determines the belief and the disbelief in the proposition, respectively. Increasing the total amount of

evidence (e) reduces the uncertainty. Note that there is a fundamental difference between an opinion where a proposition is equally provable and disprovable and one where we have complete uncertainty about the proposition. For instance, opinion $(0, 0, 1)$ indicates that there is no evidence either supporting or contradicting the proposition, i.e. $n = p = 0$, whereas opinion $(0.5, 0.5, 0)$ indicates that $n = p = \infty$.

We use the notation $p(x) \stackrel{\text{def}}{=} 2 \frac{x_b}{x_u}$ to denote the amount of supporting evidence underlying opinion x , and likewise $n(x) \stackrel{\text{def}}{=} 2 \frac{x_d}{x_u}$ for the amount of ‘negative’ evidence. Moreover, we use the notation $e(x) = p(x) + n(x)$ to represent the total amount of evidence underlying opinion x .

SL provides a number of operators to combine opinions. One of the fundamental operations is the combination of evidence from multiple sources. Consider the following scenario. Alice has evidence (p_1, n_1) about some proposition. She forms an opinion $x_1 = (p_1, n_1, 2)/(p_1 + n_1 + 2)$. Later she collects additional evidence (p_2, n_2) independent of the first evidence. Based on the second body of evidence alone, she would arrive at opinion $x_2 = (p_2, n_2, 2)/(p_2 + n_2 + 2)$. If she combines all the evidence, she obtains opinion

$$x = \frac{(p_1 + p_2, n_1 + n_2, 2)}{p_1 + p_2 + n_1 + n_2 + 2}. \quad (4)$$

The combined opinion x is expressed as a function of x_1 and x_2 via the so-called ‘consensus’ rule; this is denoted as $x = x_1 \oplus x_2$.

Definition 2 (Consensus) [15] Let $x, y \in \Omega$. The consensus $x \oplus y \in \Omega$ is defined as

$$x \oplus y \stackrel{\text{def}}{=} \frac{(x_u y_b + y_u x_b, x_u y_d + y_u x_d, x_u y_u)}{x_u + y_u - x_u y_u}. \quad (5)$$

Eq. (5) precisely corresponds to (4). Note that $x \oplus y = y \oplus x$. Furthermore, the consensus operation is associative, i.e. $x \oplus (y \oplus z) = (x \oplus y) \oplus z$. These properties are exactly what one intuitively expects from an operation that combines evidence. It is worth noting that the evidence has to be *independent* for the \oplus rule to apply. Combining dependent evidence would lead to the problem of double counting evidence. We formalize and discuss this problem in Section 3.3.

The second important operation in SL is the transfer of opinions from one party to another. Consider the following scenario. Alice has opinion x about Bob’s trustworthiness. Bob has opinion y about some proposition P . He informs Alice of his opinion. Alice now has to form an opinion about P . The standard solution to this problem is that Alice applies an x -dependent weight to Bob’s opinion y [4, 23, 26, 31, 35]. This is the so called ‘discounting’. The following formula is usually applied.

Definition 3 (Discounting) [15] Let $x, y \in \Omega$. The discounting of opinion y using opinion x is denoted as $x \otimes y \in \Omega$, and is defined as

$$x \otimes y \stackrel{\text{def}}{=} (x_b y_b, x_b y_d, x_d + x_u + x_b y_u). \quad (6)$$

It holds that $x \otimes y \neq y \otimes x$ and that $x \otimes (y \otimes z) = (x \otimes y) \otimes z$. The discounting rule (6) is *not* distributive w.r.t. consensus, i.e. $(x \oplus y) \otimes z \neq (x \otimes z) \oplus (y \otimes z)$ and $x \otimes (y \oplus z) \neq (x \otimes y) \oplus (x \otimes z)$.

In SL, a trust network can be modeled with a combination of consensus and discounting operators. The consensus operator is used to aggregate trust information from different sources, while the discounting operator is used to implement trust transitivity. Note that, in

a trust network, SL distinguishes two types of trust relationship: *functional trust*, which represents the opinion about an entity's ability to provide a specific function, and *referral trust*, which represents the opinion about an entity's ability to provide recommendations about other entities. Referral trust is assumed to be transitive, and a trust chain is said to be valid if the last edge of the chain represents functional trust and all previous edges represent referral trust.

3 Limitations of Subjective Logic

Our desideratum is a novel reputation metric that has the advantages of both SL and flow-based models. On one hand, we aim at an automated procedure for computing reputation as in flow-based approaches. On the other hand, we aim to determine the confidence in reputation values by making uncertainty explicit as in SL. In this section, we discuss the limitations of SL. Then, in Section 4 we show how these limitations affect a naïve approach that combines flow-based reputation and SL.

3.1 Dogmatic Opinions

Definition 4 The extreme points corresponding to full belief (B), full disbelief (D) and full uncertainty (U) are defined as

$$B \stackrel{\text{def}}{=} (1, 0, 0) \quad D \stackrel{\text{def}}{=} (0, 1, 0) \quad U \stackrel{\text{def}}{=} (0, 0, 1). \quad (7)$$

The special points B, D, U behave as follows regarding the consensus operation: $B \oplus x = B$; $D \oplus x = D$; $U \oplus x = x$. The full uncertainty U behaves like an additive zero.

With respect to the discounting rule, the special points B, D, U behave as $B \otimes x = x$, $D \otimes x = U$, $U \otimes x = U$, $x \otimes U = U$, $x \otimes B = (x_b, 0, 1 - x_b)$, $x \otimes D = (0, x_b, 1 - x_b)$.

Opinions that have $u = 0$ (i.e., lying on the line between B and D) are called ‘dogmatic’ opinions. They have to be treated with caution, since they have $e = \infty$ and therefore overwhelm other opinions when the consensus \oplus is applied. We will come back to this issue in Section 5.1.

3.2 Counter-intuitive behaviour of the \otimes operation

We observe that the discounting rule \otimes does not have a natural interpretation in terms of evidence handling. For instance, if we compute the positive evidence contained in $x \otimes y$ we get, using p, n notation as introduced in Section 2.2,

$$\begin{aligned} p(x \otimes y) &= 2 \frac{(x \otimes y)_b}{(x \otimes y)_u} = 2 \frac{x_b y_b}{x_d + x_u + x_b y_u} = 2 \frac{\frac{1}{4} p(x) p(y)}{\frac{1}{2} n(x)/y_u + 1/y_u + \frac{1}{2} p(x)} \\ &= \frac{p(x) p(y)}{[2 + n(x)][1 + 2p(y) + 2n(y)] + p(x)} \end{aligned} \quad (8)$$

where in the final step we have used $1/y_u = 1 + 2p(y) + 2n(y)$. Similarly, we get

$$n(x \otimes y) = 2 \frac{(x \otimes y)_d}{(x \otimes y)_u} = \frac{p(x) n(y)}{[2 + n(x)][1 + 2p(y) + 2n(y)] + p(x)}. \quad (9)$$

Eqs. (8) and (9) are complicated functions of the amounts $p(x), n(x), p(y), n(y)$. We show that they do not have a clear and well-defined interpretation in terms of evidence handling. For instance, Eq. (8) does not allow us to determine whether the evidence underlying $x \otimes y$ originates from x or y ; one can argue that $p(x \otimes y)$ is either a “fraction” of $p(y)$, a “fraction” of $p(x)$, or contains evidence from both y and x . Even if we try to interpret in the standard way (see Section 2.2), we can observe that the factor multiplying $p(y)$ in (8) and $n(y)$ in (9) depends on $p(y)$ and $n(y)$. Hence, the evidence underlying $x \otimes y$ is not simply an x -dependent multiple of the evidence underlying y . In fact, the examples below show that in some cases the contribution from y completely disappears from the equation.

Example 2

Let $x, y \in \Omega$, with $n(x) = 0, n(y) = 0$ and $p(y) \gg p(x)$. Then $p(x \otimes y) \approx p(x)/4$.

Example 3

Let $x, y \in \Omega$, with $n(x) = 0, p(y) = 0$ and $n(y) \gg p(x)$. Then $n(x \otimes y) \approx p(x)/4$.

In both these examples, y is based on a lot of evidence; but even if x contains a lot of belief, none of y ’s evidence survives in $x \otimes y$. We conclude that the discounting operation \otimes gives counter-intuitive results.

The \otimes rule is inspired by a *probabilistic* interpretation of opinions. The probabilistic interpretation might suggest that it is natural to multiply probabilities, i.e. that the expressions $(x \otimes y)_b = x_b y_b$ and $(x \otimes y)_d = x_b y_d$ are intuitively correct. However, we argue that this is not at all self-evident. When discounting y through x , the uncertainties in x induce an x -dependent probability distribution on y . This can be thought of as an additional layer of uncertainty about beta distributions. Let (3) describe opinion y ; then the discounting through x introduces uncertainty about the parameters p and n in the equation (a probability distribution on p and n). It is not at all self-evident that the resulting opinion is $x \otimes y$ as prescribed by Def. 3. It would make equal sense to replace the discounting factor x_b by e.g. the expectation $x_b + a x_u$. In this paper we do not pursue such an approach based on distributions, but we mention it in order to point out that the \otimes rule is not necessarily well-founded.

The fact that SL employs on the one hand a consensus rule based on *adding evidence* and on the other hand a discounting rule based on *multiplying opinions* leads to a lack of ‘cooperation’ between the \oplus and \otimes operations. Most importantly, the \otimes is not distributive with respect to \oplus ,

$$(x \otimes y) \oplus (x \otimes z) \neq x \otimes (y \oplus z). \quad (10)$$

Consider the following scenario.

Example 4 Alice has opinion x about Bob’s trustworthiness in providing recommendations. Bob experiments with chocolate on Monday and forms an opinion y about its medicinal qualities. On Tuesday he does some more of the same kind of experiments and forms an independent opinion z . He informs Alice of y, z and his final opinion $y \oplus z$. What should Alice think about the medicinal qualities of chocolate? One approach is to say that Alice should appraise opinions y and z separately, yielding $(x \otimes y) \oplus (x \otimes z)$ (Note that the two occurrences of x represent the very same opinion, i.e. the evidence underlying the two occurrences is the same). Another approach is to weight Bob’s combined opinion, yielding $x \otimes (y \oplus z)$. Intuitively the two approaches should yield exactly the same opinion, yet the SL rules give (10).

We now present a numerical example that illustrates the issue discussed above.




Fig. 1: Trust network representing the scenario in Example 4. A is Alice, B_1 and B_2 are Bob on Monday and Tuesday respectively, and P is the proposition “chocolate has medicinal qualities”. Referral trust is drawn as a full line, and functional trust as a double full line.

Example 5 Figure 1 shows the trust network representing the scenario in Example 4. To highlight that opinions x and y are independent, in the figure we abuse the network notation and duplicate the node representing Bob: B_1 represents Bob on Monday and B_2 represents Bob on Tuesday. The edge between Alice (A) and Bob (dashed rectangle) represents Alice’s opinion x about Bob’s recommendations. This opinion concerns Bob’s recommendations regardless of when they are formed (e.g., on Monday or on Tuesday).

Suppose that Alice’s opinion about Bob’s trustworthiness is $x = (0.6, 0.1, 0.3)$, and Bob’s opinions about the proposition P are $y = (0.3, 0.6, 0.1)$ and $z = (0.5, 0.2, 0.3)$. We are interested in Alice’s opinion w about P based on Bob’s recommendations. As discussed in the previous example, we have two approaches to compute such an opinion:

1. $w = (x \otimes y) \oplus (x \otimes z) = (0.314, 0.341, 0.345)$
2. $w = x \otimes (y \oplus z) = (0.227, 0.324, 0.449)$

Clearly, the two approaches yield different opinions, contradicting the intuitive expectation.

3.3 Double counting of evidence in Trust Networks

The \oplus rule imposes constraints on the evidence that can be aggregated: it requires evidence to be independent [15]. In the literature, however, there is no well-defined notion of evidence independence. Some researchers [33, 45] assume that pieces of evidence are independent if they are obtained from independent sources, where two sources are said to be independent if they measure completely unrelated features. This definition, however, is too restrictive. For instance, the evidence collected by a sensor at different points in time can also be independent.

Evidence is usually extracted from “observations” of a system. In this work we adopt a notion of evidence independence based on the independence of observations, which we define in the same way as independence of random variables.

Definition 5 (Independent Evidence) Let o_i and o_j be two observations and e_i and e_j the evidence obtained from o_i and o_j respectively. We say that o_i and o_j are independent if and only if $\text{Prob}(o_i|o_j) = \text{Prob}(o_i)$. Moreover, we say that e_i and e_j are independent if they are obtained from independent observations.

Intuitively, observation independence requires that the probability that o_i happens, assuming that o_j happened, is the same as the probability that o_i happens regardless of o_j . The definition above can be extended to opinions:

Definition 6 (Independent Opinions) Let $x, y \in \Omega$ be opinions. We say that x and y are independent if and only if the evidence underlying x and y is independent.

Combining dependent evidence leads to the problem of double counting evidence.

Definition 7 (Double Counting) Let $x, y \in \Omega$. In an expression of the form $x \oplus y$ we say that there is double counting if there is dependence between the evidence underlying x and the evidence underlying y .

Intuitively, dependent evidence shares “part” of the evidence. Therefore, aggregating dependent evidence leads to counting some part of evidence more than once.

Example 6 Consider the expression $(y \otimes x) \oplus (z \otimes x)$, where both occurrences of x are obtained from the same observation. The evidence underlying x is contained on the left side as well as the right side of the ‘ \oplus ’. This is a clear case of double counting.

Example 7 Consider the expression $(x \otimes y) \oplus (x \otimes z)$, again with both instances ‘ x ’ coming from the same observation. The evidence underlying x is contained on the left side as well as the right side of the ‘ \oplus ’, but less evidently than in Example 6, because now x is used for discounting. In Section 3.2 we showed that the \otimes rule causes evidence from x to end up in $x \otimes y$ in a complicated way. Hence the opinions $x \otimes y$ and $x \otimes z$ are *not* independent, which causes double counting of x in the expression $(x \otimes y) \oplus (x \otimes z)$.

It is worth noting that double counting of x in the expression $(x \otimes y) \oplus (x \otimes z)$ can also be observed in Example 5. Indeed, the uncertainty in $(x \otimes y) \oplus (x \otimes z)$ is lower than the uncertainty in $x \otimes (y \oplus z)$, indicating that the result contains more evidence when the trust network is represented using the first expression.

To avoid the problem of double counting, SL requires that the trust network is expressed in a *canonical form* [20, 22], where all trust paths are independent. Intuitively, a trust network expression is in canonical form if every edge appears only once in the expression.

Example 8 Consider the two trust network expressions representing the trust network in Figure 1 given in Example 4: $(x \otimes y) \oplus (x \otimes z)$ and $x \otimes (y \oplus z)$. The first expression is not in canonical form as opinion x occurs twice in the expression; the second expression is in canonical form as every edge appears only once in the expression. Thus, $x \otimes (y \oplus z)$ is the proper representation of the trust network in Figure 1.

In the next section we show that it is not always possible to express a trust network in canonical form. As suggested in [20, 22], this issue can be addressed by removing some edges from the network. This, however, means discarding part of the trust information, thus reducing the quality of reputation values.

4 Combining flow-based reputation and Subjective Logic

This section presents a naïve approach that combines flow-based reputation and SL. We illustrate the limitations of such a naïve approach.

We first introduce some notation and definitions. Flow-based reputation models usually assume that users who are honest during transactions are also honest in reporting their ratings [23]. This assumption, however, does not hold in many real-life situations [1]. Thus, as it is done in SL, we distinguish between referral trust and functional trust (see Section 2). We use A to represent direct referral trust and T to represent direct functional trust. R denotes the final referral trust and F the final functional trust.




Fig. 2: Example of a trust network that is problematic for Subjective Logic.

Definition 8 For n users, the *direct referral trust matrix* A is an $n \times n$ matrix, where $A_{xy} \in \Omega$ (with $x \neq y$) is the direct referral trust that user x has in user y , and $A_{xx} = (0, 0, 1)$ for all x .

Note that we impose the condition $A_{xx} = (0, 0, 1)$ in order to prevent artifacts caused by self-rating [35].

Let T_{jP} be the opinion of user j about some proposition P , and let R_{ij} be i 's (possibly indirect) opinion about the trustworthiness of user j . The opinion of user i about P based on direct and indirect evidence can be computed using the following equation:

$$F_{iP} = T_{iP} \oplus \bigoplus_{j:j \neq i} (R_{ij} \otimes T_{jP}). \quad (11)$$

Eq. (11) computes the final functional trust F_{iP} by combining user i 's direct opinion T_{iP} with all the available opinions of other users, $\{T_{jP}\}_{j \neq i}$. The opinion received from j is weighted with the 'reputation' R_{ij} .

To find R_{ij} , we could try a recursive approach¹ inspired by Eq. (1):

$$R_{ij} = A_{ij} \oplus \bigoplus_{k:k \neq i} (R_{ik} \otimes A_{kj}) \quad \text{for } i \neq j. \quad (12)$$

To demonstrate the problems that occur in this naïve combination of flow-based reputation and SL, we consider the trust networks shown in Figures 2 and 3. Referral trust is drawn as a full line, functional trust as a double full line.

Figure 2 is a variant of a network discussed in [20]. We are interested in determining the opinion F_{1P} of user 1 about some proposition P . User 1 does not have any direct evidence. The only functional trust about P comes from user 7. Thus, we have to determine R_{17} , user 1's referral trust in user 7. This is done by taking A_{67} with the proper weight, namely R_{16} . Continuing this recursive approach using (12) gives

$$F_{1P} = [A_{12} \otimes A_{23} \otimes A_{34} \otimes A_{46} \oplus (A_{12} \otimes A_{23} \otimes A_{34} \otimes A_{45} \oplus A_{12} \otimes A_{23} \otimes A_{35}) \otimes A_{56}] \otimes A_{67} \otimes T_{7P}. \quad (13)$$

This, however, is a problematic result. Recall that SL requires trust networks to be expressed in a canonical form. If this restriction is not satisfied, we face the problem of 'double counting' opinions, i.e. applying the \oplus operation to opinions that are not independent (Def. 7).

In Figure 2, consider the case $A_{45} = U$. The canonical solution for this case is

$$F_{1P} = A_{12} \otimes A_{23} \otimes (A_{34} \otimes A_{46} \oplus A_{35} \otimes A_{56}) \otimes A_{67} \otimes T_{7P} \quad (14)$$

¹ We consider only scenarios where all entities *publish* their opinions. If opinions are not published but communicated solely over the network links, then a recursive equation containing $A \otimes R$ (the opposite order of $R \otimes A$) applies.




Fig. 3: Example of a trust network containing a loop.

whereas (13) yields

$$F_{1P} = (A_{12} \otimes A_{23} \otimes A_{34} \otimes A_{46} \oplus A_{12} \otimes A_{23} \otimes A_{35} \otimes A_{56}) \otimes A_{67} \otimes T_{7P}. \quad (15)$$

In (15) the A_{12} and A_{23} are double-counted. We conclude that the naïve recursive equation (12) does not properly reproduce the canonical solution.

There is a further problem, unrelated to the naïve recursive approach. As was shown in [20], it is not even possible to transform the trust network in Figure 2 into a canonical form in the general case $A_{45} \neq U$.

The problems become even worse when the trust network contains loops, e.g. a loop as shown in Figure 3. Here too, there is no canonical form. Applying the recursive approach to Figure 3 gives $R_{13} = R_{12} \otimes A_{23}$, with $R_{12} = A_{12} \oplus R_{12} \otimes A_{23} \otimes A_{32}$. Repeatedly substituting the latter into itself yields

$$R_{12} = A_{12} \oplus [A_{12} \oplus \{A_{12} \oplus \dots\} \otimes A_{23} \otimes A_{32}] \otimes A_{23} \otimes A_{32}. \quad (16)$$

We observe that taking opinion A_{32} into account causes excessive double-counting of A_{12} . If the loop is somehow discarded, then the information contained in A_{32} is destroyed.

In conclusion, (i) generic trust networks with several connections $A_{ij} \neq U$ cannot be handled with SL because there is no canonical form for them that avoids double-counting; (ii) even when there is a canonical result, this result cannot be reproduced by a straightforward recursive approach.

5 Subjective Logic revisited

This section presents a new, fully evidence-based approach to SL. We refer to the resulting opinion algebra as Evidence-Based Subjective Logic or EBSL.

5.1 Excluding dogmatic opinions

As mentioned in Section 3.1, dogmatic opinions are problematic when the \oplus operation is applied to them. Furthermore, a dogmatic opinion corresponds to an infinite amount of evidence, which in our context is not realistic. In the remainder of this paper, we will exclude dogmatic opinions. We will work with a reduced opinion space defined as follows.

Definition 9 The *opinion space excluding dogmatic opinions* is denoted as Ω' and is defined as $\Omega' \stackrel{\text{def}}{=} \{(b, d, u) \in [0, 1) \times [0, 1) \times (0, 1] \mid b + d + u = 1\}$.

We are by no means the first to do this; in fact, the exclusion of dogmatic opinions was proposed as an option in the very early literature on SL [22].

5.2 The relation between evidence and opinions: a simplified justification

We make a short observation about the mapping between evidence and opinions. As was mentioned in Section 2.2, there is a one-to-one mapping (2) based on the analysis of probability distributions (Beta distributions). Here we show that there is a shortcut: the same mapping can also be obtained in a much simpler way, based on constraints.

Theorem 1 *Let $p \geq 0$ be the amount of evidence that supports ‘belief’; let $n \geq 0$ be the amount of evidence that supports ‘disbelief’. Let $x = (b, d, u) \in \Omega'$ be the opinion based on the evidence. If we demand the following four properties,*

1. $b/d = p/n$
2. $b + d + u = 1$
3. $p + n = 0 \Rightarrow u = 1$
4. $p + n \rightarrow \infty \Rightarrow u \rightarrow 0$

then the relation between x and (p, n) has to be

$$x = (b, d, u) = \frac{(p, n, c)}{p + n + c}; \quad (p, n) = c \frac{(b, d)}{u} \quad (17)$$

where $c > 0$ is a constant.

Proof: Property 1 gives $(b, d) = (p, n)/f(p, n)$ where f is some function. Combined with property 2 we then have $\frac{p+n}{f(p, n)} = 1 - u$. Property 3 then gives $\frac{0}{f(0,0)} = 0$, while property 4 gives $\lim_{p+n \rightarrow \infty} \frac{p+n}{f(p, n)} = 1$. The latter yields $f(p, n) = p + n + c$ where c is some constant. Allowing $c < 0$ would open the possibility of components of x being negative. Thus we must have $c \geq 0$. The requirement $\frac{0}{f(0,0)} = 0$ eliminates the possibility of having $c = 0$, since $c = 0$ would yield division by zero. Finally, setting $u = c/(p + n + c)$ is necessary to satisfy property 2. \square

Theorem 1 shows that we can derive a formula similar to (2), based on minimal requirements which make intuitive sense. Only the constant c is not fixed by the imposed constraints; it has to be determined from the context. One can interpret c as a kind of soft threshold on the amount of evidence: beyond this threshold one starts gaining enough confidence from the evidence to form an opinion.

We observe that (17) with its generic constant c is already sufficient to derive the consensus rule \oplus , i.e. *the consensus rule does not require $c = 2$.*

Lemma 1 *The mapping (17) with arbitrary c implies the consensus rule \oplus as specified in Def. 2.*

Proof: Consider $x = (b_1, d_1, u_1) = (p_1, n_1, c)/(p_1 + n_1 + c)$ and $y = (b_2, d_2, u_2) = (p_2, n_2, c)/(p_2 + n_2 + c)$. An opinion formed from the combined evidence $(p_1 + p_2, n_1 + n_2)$ according to (17) is given by $(b, d, u) = \frac{(p_1 + p_2, n_1 + n_2, c)}{p_1 + n_1 + p_2 + n_2 + c}$. Substituting $p_i = c \frac{b_i}{u_i}$ and $n_i = c \frac{d_i}{u_i}$ yields, after some simplification, $(b, d, u) = \frac{(u_1 b_2 + u_2 b_1, u_1 d_2 + u_2 d_1, u_1 u_2)}{u_1 + u_2 - u_1 u_2}$. \square

Theorem 1 and Lemma 1 improve our understanding of evidence-based opinion forming and of the consensus rule.

Furthermore, in (2) and (3) we can replace ‘2’ by c and the expectation of t , obtained by integrating the Beta distribution times t , is still $x_b + ax_u$! Therefore, in the remainder of the paper, we will work with a re-defined version of the $p(x)$ and $n(x)$ functions (Section 2.2). The new version has a general value $c > 0$ instead of $c = 2$.

Definition 10 Let $x = (x_b, x_d, x_u) \in \Omega'$. We define the notation $p(x)$, $n(x)$ and $e(x)$ as

$$p(x) \stackrel{\text{def}}{=} c \frac{x_b}{x_u}; \quad n(x) \stackrel{\text{def}}{=} c \frac{x_d}{x_u}; \quad e(x) \stackrel{\text{def}}{=} p(x) + n(x). \quad (18)$$

5.3 Scalar multiplication

Our next contribution has more impact. We define an operation on opinions that is equivalent to a scalar multiplication on the total amount of evidence.

Definition 11 (Scalar multiplication) Let $x = (x_b, x_d, x_u) \in \Omega'$ and let $\alpha \geq 0$ be a scalar. We define the product $\alpha \cdot x$ as

$$\alpha \cdot x \stackrel{\text{def}}{=} \frac{(\alpha x_b, \alpha x_d, x_u)}{\alpha(x_b + x_d) + x_u}. \quad (19)$$

Lemma 2 Let $x \in \Omega'$ and $\alpha \geq 0$. The scalar multiplication as specified in Def. 11 has the following properties:

1. $\alpha \cdot x \in \Omega'$.
2. $0 \cdot x = U$.
3. $1 \cdot x = x$.
4. For $n \in \mathbb{N}$, $n \geq 2$, it holds that $n \cdot x = \underbrace{x \oplus x \oplus \dots \oplus x}_n$.
5. The evidence underlying $\alpha \cdot x$ is α times the evidence underlying x , i.e. $p(\alpha \cdot x) = \alpha p(x)$ and $n(\alpha \cdot x) = \alpha n(x)$.
6. If $\alpha \neq 0$ then $(\alpha \cdot x)_b / (\alpha \cdot x)_d = x_b / x_d$.

Proof: Property 1: It is readily verified that the components add to 1, using $x_b + x_d = 1 - x_u$. Since $\alpha \geq 0$, all three components of $\alpha \cdot x$ are nonnegative, since $x \in \Omega'$. Properties 2 and 3: Found directly by substituting $\alpha = 0$ resp. $\alpha = 1$ in (19). Property 4: Consider $n = 2$. Setting $\alpha = 2$ in (19) yields $2 \cdot x = \frac{(2x_b, 2x_d, x_u)}{2 - x_u} = \frac{(2x_u x_b, 2x_u x_d, x_u^2)}{2x_u - x_u^2} = x \oplus x$. The rest follows by induction. Property 5: We use (17) to map between opinions and evidence. The positive evidence of x is $p(x) = cx_b/x_u$. The positive evidence of $\alpha \cdot x$ is $p(\alpha \cdot x) = c \frac{\alpha x_b}{\alpha(1-x_u)+x_u} / \frac{x_u}{\alpha(1-x_u)+x_u} = \alpha cx_b/x_u = \alpha p(x)$. The proof for $n(\alpha \cdot x)$ is analogous. Property 6: Follows directly by dividing the first and second component of $\alpha \cdot x$. \square

Lemma 3 (Distributivity of the scalar multiplication) Let $x, y \in \Omega'$. Let $\alpha, \beta \geq 0$. Then it holds that

$$\alpha \cdot (x \oplus y) = (\alpha \cdot x) \oplus (\alpha \cdot y) \quad \text{and} \quad (\alpha + \beta) \cdot x = (\alpha \cdot x) \oplus (\beta \cdot x). \quad (20)$$

Proof: We make extensive use of $\alpha \cdot x \propto (\alpha x_b, \alpha x_d, x_u)$. First part: On the one hand, $\alpha \cdot (x \oplus y) \propto (\alpha[x_u y_b + y_u x_b], \alpha[x_u y_d + y_u x_d], x_u y_u)$. We have not written the normalization factor; it is not necessary since we already know that the result is normalized. On the other hand, $(\alpha \cdot x) \oplus (\alpha \cdot y) \propto (x_u[\alpha y_b] + y_u[\alpha x_b], x_u[\alpha y_d] + y_u[\alpha x_d], x_u y_u)$.

Second part: On the one hand, $(\alpha + \beta) \cdot x \propto ([\alpha + \beta]x_b, [\alpha + \beta]x_d, x_u)$. On the other hand, $(\alpha \cdot x) \oplus (\beta \cdot x) \propto ([\beta \cdot x]_u[\alpha \cdot x]_b + [\alpha \cdot x]_u[\beta \cdot x]_b, [\beta \cdot x]_u[\alpha \cdot x]_d + [\alpha \cdot x]_u[\beta \cdot x]_d, [\alpha \cdot x]_u[\beta \cdot x]_u) \propto (x_u[\alpha x_b] + x_u[\beta x_b], x_u[\alpha x_d] + x_u[\beta x_d], x_u^2) \propto (\alpha x_b + \beta x_b, \alpha x_d + \beta x_d, x_u)$. \square

Lemma 4 Let $x \in \Omega'$ and $\alpha, \beta \geq 0$. Then $\alpha \cdot (\beta \cdot x) = (\alpha\beta) \cdot x$.

Proof: $\alpha \cdot (\beta \cdot x) \propto \alpha \cdot (\beta x_b, \beta x_d, x_u) \propto (\alpha\beta x_b, \alpha\beta x_d, x_u)$. \square

5.4 New discounting rule

We propose a new approach to discounting: instead of multiplying (part of) the *opinions* we multiply the *evidence*. The multiplication is done using our scalar multiplication rule (Def. 11). We return to the example where Alice has an opinion $x \in \Omega'$ about the trustworthiness of Bob, and Bob has an opinion $y \in \Omega'$ about some proposition P . We propose a discounting of the form $g(x) \cdot y$, where $g(x) \geq 0$ is a scalar that indicates which fraction of Bob's evidence is accepted by Alice. One can visualize the discounting as a physical transfer of evidence from Bob to Alice, during which only a fraction $g(x)$ survives, due to Alice's mistrust and/or uncertainty. It is desirable to set $g(x)$ in the range $[0, 1]$: allowing $g(x) < 0$ would lead to negative *amounts* of evidence (not to be confused with the term 'negative evidence' which is used for evidence that contradicts the proposition P); allowing $g(x) > 1$ would "amplify" evidence, i.e., create new evidence out of nothing, which is clearly unrealistic.

It makes intuitive sense to set $\lim_{x \rightarrow B} g(x) = 1$, $\lim_{x \rightarrow D} g(x) = 0$ and $g(U) = 0$, or even to set $g(x) = \tilde{g}(x_b)$, i.e. a function of x_b only, with $\tilde{g}(0) = 0$ and $\tilde{g}(1) = 1$. For instance, we could set $g(x) = x_b$.² On the other hand, it could also make sense to set $g(U) > 0$, which would represent the "benefit of the doubt". An intuitive choice would then be $g(x) = x_b + ax_u$, i.e. the expectation value corresponding to x . We postpone the precise details of how the function g can/should be chosen, and introduce a very broad definition.

Definition 12 (New generic discounting rule \boxtimes) Let $x, y \in \Omega'$. Let $g : \Omega' \rightarrow [0, 1]$ be a function. We define the operation \boxtimes as

$$x \boxtimes y \stackrel{\text{def}}{=} g(x) \cdot y = \frac{(g(x)y_b, g(x)y_d, y_u)}{(y_b + y_d)g(x) + y_u} \quad (21)$$

with the \cdot operation as specified in Def. 11.

Differently from \otimes , the operator \boxtimes has a well-defined interpretation in terms of evidence handling. The following theorem states that the evidence underlying $x \boxtimes y$ is a fraction of the evidence underlying y defined by a scalar weight depending on x .

Theorem 2 Let $x, y \in \Omega'$. The operation \boxtimes (Def. 12) has the following properties:

1. $x \boxtimes y \in \Omega'$.
2. $p(x \boxtimes y) = g(x)p(y)$ and $n(x \boxtimes y) = g(x)n(y)$.
3. $(x \boxtimes y)_b / (x \boxtimes y)_d = y_b / y_d$.
4. $x \boxtimes U = U$.
5. Discounting cannot decrease uncertainty, i.e. $(x \boxtimes y)_u \geq y_u$.

Proof: Property 1 follows from $x \boxtimes y = g(x) \cdot y$ and the first property in Lemma 2. Property 2: We compute $p(x \boxtimes y) = (x \boxtimes y)_b / (x \boxtimes y)_u$ using the definition (21), which yields $g(x)y_b / y_u = g(x)p(y)$. For $n(x \boxtimes y)$ the derivation is analogous. Property 3: Follows directly by dividing the belief and disbelief part of (21). Property 4: Follows by setting $y_b = 0$, $y_d = 0$ and $y_u = 1$ in (21). Property 5: We have $(x \boxtimes y)_u = \frac{y_u}{(y_b + y_d)g(x) + y_u}$. Since $y_b + y_d + y_u = 1$ and $g(x) \in [0, 1]$, the denominator of the fraction lies in the range $[y_u, 1]$. \square

² Even though taking $g(x) = x_b$ appears to be similar to the \otimes rule (6), we show in Section 7 that they have a very different behavior.

Corollary 1 Let $x, y \in \Omega'$ with $g(x) > 0$. Let $\lim_{y \rightarrow B} g(y) = 1$, $\lim_{y \rightarrow D} g(y) = 0$ and $g(U) = 0$. Then the extreme points B, D, U have the following behavior with respect to the new discounting rule \boxtimes ,

$$\begin{aligned} \lim_{y \rightarrow B} y \boxtimes x &= x & \lim_{y \rightarrow B} x \boxtimes y &= B \\ \lim_{y \rightarrow D} y \boxtimes x &= U & \lim_{y \rightarrow D} x \boxtimes y &= D \\ U \boxtimes x &= U. \end{aligned} \quad (22)$$

We stress again that the whole ‘dogmatic’ line between B and D is not part of the opinion space Ω' , so that we avoid having to deal with infinite amounts of evidence.

Theorem 3 There is no function $g : \Omega' \rightarrow [0, 1]$ such that $x \boxtimes y = x \otimes y$ for all $x, y \in \Omega'$.

Proof: On the one hand, we have $x \otimes y = (x_b y_b, x_b y_d, 1 - x_b y_b - x_b y_d)$. On the other hand, $x \boxtimes y = \frac{(g(x)y_b, g(x)y_d, y_u)}{g(x)(y_b + y_d) + y_u}$. Demanding that they are equal yields, after some rewriting, $g(x) = x_b y_u / [1 - x_b(1 - y_u)]$. This requires $g(x)$, which is a function of x only, to be a function of y_u as well. \square

Being based on the scalar multiplication rule (and hence ultimately on the \oplus rule), our operation \boxtimes has several properties that \otimes lacks: (i) right-distribution; (ii) permutation symmetry of parties that transfer evidence. This is demonstrated below.

Lemma 5 Let $x, y, z \in \Omega'$. Then

$$x \boxtimes (y \oplus z) = (x \boxtimes y) \oplus (x \boxtimes z). \quad (23)$$

Proof: It follows trivially from $x \boxtimes (y \oplus z) = g(x) \cdot (y \oplus z)$ and Lemma 3. \square

This distributive property resolves the issue discussed in Section 3.2: using the \boxtimes operator, it does not matter if y and z are combined before or after the discounting. This solves the inconsistency caused by the \otimes operation.

Notice also that the left-hand side of (23) obviously is not double-counting x ; hence also the expression on the right-hand side does not double-count x . In contrast, the right-hand side expression with \otimes instead of \boxtimes would be double-counting. We come back to this point in Section 5.6.

Lemma 6 Let $y, x_1, x_2 \in \Omega'$. Then

$$x_1 \boxtimes (x_2 \boxtimes y) = x_2 \boxtimes (x_1 \boxtimes y). \quad (24)$$

Proof: $x_1 \boxtimes (x_2 \boxtimes y) = g(x_1) \cdot (g(x_2) \cdot y)$. Using Lemma 4 this reduces to $(g(x_1)g(x_2)) \cdot y$. Exactly the same reduction applies to $x_2 \boxtimes (x_1 \boxtimes y)$. \square

Lemma 6 generalizes to chains of discounting:

$$x_1 \boxtimes (x_2 \boxtimes (\cdots (x_N \boxtimes y))) = \left[\prod_{i=1}^N g(x_i) \right] \cdot y. \quad (25)$$

Expression (25) is invariant under permutation of the opinions x_1, \dots, x_N .

Note that one property of the \otimes rule is *not* generically present in the \boxtimes rule: associativity. Whereas the old rule has $x \otimes (y \otimes z) = (x \otimes y) \otimes z$, the new rule has

$$\begin{aligned} x \boxtimes (y \boxtimes z) &= [g(x)g(y)] \cdot z \quad \text{versus} \\ (x \boxtimes y) \boxtimes z &= (g(x \boxtimes y)) \cdot z = [g(g(x) \cdot y)] \cdot z. \end{aligned} \quad (26)$$

However, it is important to realize that *the lack of associativity is not a problem*. The transfer of evidence along a chain has a very clear ordering, which determines the order in which the \boxtimes operations have to be performed. (See Section 6.)

Also note that \boxtimes does not have a left-distribution property for arbitrarily chosen g . It takes some effort to define a reasonable function g that yields left-distributivity.

Lemma 7 *There is no function $g : \Omega' \rightarrow [0, 1]$ that satisfies $\lim_{s \rightarrow B} g(s) = 1$ and gives $(x \oplus y) \boxtimes z = (x \boxtimes z) \oplus (y \boxtimes z)$ for all $x, y, z \in \Omega'$.*

Proof: We consider the limit $x \rightarrow B, y \rightarrow B$. On the one hand, $(x \oplus y) \boxtimes z \rightarrow B \boxtimes z = z$. On the other hand, $(x \boxtimes z) \oplus (y \boxtimes z) \rightarrow z \oplus z$. \square

It may look surprising that we cannot achieve left-distributivity with a function g chosen from a very large function space with only a single constraint. (And a very reasonable-looking constraint at that). But left-distributivity requires $g(x \oplus y) = g(x) + g(y)$, which conflicts with the constraint $\lim_{s \rightarrow B} g(s) = 1$.

5.5 New specific discounting rule

One way to satisfy $g(x \oplus y) = g(x) + g(y)$ is by setting $g(x) \propto p(x)$. This approach, however, causes some complications. Suppose we define $g(x) = p(x)/\theta$, where θ is some constant. If the amount of positive evidence ever exceeds θ , then the discounting factor becomes larger than 1, i.e. amplification instead of reduction, which is an undesirable property. If we redefine g such that factors larger than 1 are mapped back to 1, then we lose the distribution property. We conclude that the “ g proportional to evidence” approach can only work if the maximum achievable amount of positive evidence in a given trust network can be upper-bounded by θ .

Definition 13 (New specific discounting rule \odot) Let $x, y \in \Omega'$. Let θ be a threshold larger than the maximum amount of positive evidence in any opinion that is used for discounting. We define the operation \odot as

$$x \odot y \stackrel{\text{def}}{=} \left(\frac{c}{\theta} \frac{x_b}{x_u} \right) \cdot y = \frac{p(x)}{\theta} \cdot y \quad (27)$$

with the \cdot operation as specified in Def. 11.

We stress again that θ depends on the interactions between entities within the system, i.e. on the structure of the trust network and the maximum amount of positive evidence in the network.

Lemma 8 (Left-distributivity of \odot) Let $x, y, z \in \Omega'$. Then

$$(x \oplus y) \odot z = (x \odot z) \oplus (y \odot z). \quad (28)$$

Proof: The right-hand side evaluates to $(x \odot z) \oplus (y \odot z) = [g(x) \cdot z] \oplus [g(y) \cdot z] = [g(x) + g(y)] \cdot z$. The g -function in Def. 13 equals positive evidence divided by θ ; hence $g(x) + g(y) = [p(x) + p(y)]/\theta$. The left-hand side evaluates to $g(x \oplus y) \cdot z$ with $g(x \oplus y) = \theta^{-1} p(x \oplus y) = [p(x) + p(y)]/\theta$. \square

Lemma 9 (Associativity of \odot) Let $x, y, z \in \Omega'$. Then

$$x \odot (y \odot z) = (x \odot y) \odot z. \quad (29)$$

Proof: From (26) we see that $x \odot (y \odot z) = [g(x)g(y)] \cdot z$, but now we have $g(x)g(y) = \theta^{-2}p(x)p(y)$. From (26) we also see $(x \odot y) \odot z = g(g(x) \cdot y) \cdot z$, but now we have $g(g(x) \cdot y) = \theta^{-1}p(\theta^{-1}p(x) \cdot y) = \theta^{-2}p(x)p(y)$. \square

We are not claiming that \odot is the proper discounting operation to use. It has the unpleasant property that the negative evidence underlying x is completely ignored in the computation of $x \odot y$. A quick-fix of the form $g(x) \propto p(x) - n(x)$ does not work since it can cause $g(x) < 0$ and therefore $x \boxtimes y \notin \Omega'$.

We note that there is no alternative g -function to the ones discussed above if linearity of g is required. This is formalized in the following lemma.

Lemma 10 *The property $g(x \oplus y) = g(x) + g(y)$ can only be achieved by setting $g(x) = \alpha p(x) + \beta n(x)$, where α and β are scalars.*

The proof is given in the Appendix. Note that, for sufficiently smooth g , it is possible to prove that the property $g(x \oplus y) = g(x) + g(y)$ implies $g(g(x) \cdot y) = g(x)g(y)$, i.e. associativity.

5.6 The new discounting rule avoids double-counting

In Section 2.2 we saw that the consensus operation \oplus should be applied only to opinions that are derived from *independent* evidence. If this restriction is not obeyed then we speak of double-counting. Example 7 showed that in the SL expression $(x \otimes y) \oplus (x \otimes z)$, the evidence in x is double-counted. If we look at the equivalent expression in which \otimes is replaced with \boxtimes , we get

$$(x \boxtimes y) \oplus (x \boxtimes z) = x \boxtimes (y \oplus z). \quad (30)$$

Here there is obviously no double-counting. Next we look at more complicated EBSL expressions.

Lemma 11 *Let $x, y \in \Omega'$ be independent opinions. Let $Q, W \in \Omega'$ be opinions independent of x and y , but with mutual dependence. Then, the expression*

$$(Q \boxtimes x) \oplus (W \boxtimes y) \quad (31)$$

does not double-count any of the evidence underlying Q and W .

Proof: The evidence underlying $Q \boxtimes x$ is the evidence from x , scalar-multiplied by $g(Q)$. Likewise, the evidence underlying $W \boxtimes y$ is a scalar multiple of $(p(y), n(y))$. Since x and y are independent, the evidence on the left and right side of the ' \oplus ' in (31) is independent. \square

Based on the result above we can conclude that:

Corollary 2 *Transporting different pieces of evidence over the same link x with the \boxtimes operation is not double-counting x .*

Thus, many expressions that are problematic in SL become perfectly acceptable in EBSL, simply because \boxtimes is just an (attenuating) evidence transport operation, whereas SL's \otimes is a very complicated thing that mixes evidence from its left and right operand (Eqs. 8 and 9).

6 Flow-based reputation with uncertainty

In this section we will use the discounting rule \boxtimes without specifying the function g . We show that EBSL can be applied to arbitrarily connected trust networks and that the simple recursive approach (12), with \otimes replaced by \boxtimes , yields consistent results that avoid the double-counting problem.

6.1 Recursive solutions using EBSL

We show that the trust networks discussed in Section 3, which are problematic in SL, can be handled in EBSL. We take the EBSL equivalent of the recursive approach in SL (11), (12), namely

$$\begin{aligned} F_{iP} &= T_{iP} \oplus \bigoplus_{j:j \neq i} (R_{ij} \boxtimes T_{jP}) \\ R_{ij} &= A_{ij} \oplus \bigoplus_{k:k \neq i} (R_{ik} \boxtimes A_{kj}) \quad \text{for } i \neq j, \end{aligned} \quad (32)$$

and demonstrate that these equations yield acceptable results in the case of the trust networks depicted in Figures 2 and 3.

For Figure 2 we obtain the EBSL equivalent of (13) by using (32) recursively as follows,

$$\begin{aligned} F_{1P} &= R_{17} \boxtimes T_{7P} \\ R_{17} &= R_{16} \boxtimes A_{67} \\ R_{16} &= (R_{14} \boxtimes A_{46}) \oplus (R_{15} \boxtimes A_{56}) \\ R_{15} &= (R_{14} \boxtimes A_{45}) \oplus (R_{13} \boxtimes A_{35}) \\ R_{14} &= R_{13} \boxtimes A_{34} \\ R_{13} &= R_{12} \boxtimes A_{23} \\ R_{12} &= A_{12}. \end{aligned}$$

By substituting these expressions from bottom to top we get the end result for F_{1P} . The result is very similar to (13), but now we have lots of brackets because \boxtimes is not associative. We inspect R_{16} ,

$$\begin{aligned} R_{16} &= ((A_{12} \boxtimes A_{23}) \boxtimes A_{34}) \boxtimes A_{46} \oplus \\ &\quad \left\{ ((A_{12} \boxtimes A_{23}) \boxtimes A_{34}) \boxtimes A_{45} \oplus (A_{12} \boxtimes A_{23}) \boxtimes A_{35} \right\} \boxtimes A_{56}. \end{aligned} \quad (33)$$

We observe that the links A_{12} and A_{23} occur three times. The link A_{34} occurs twice. However, the computation of R_{16} only relies on the evidence in A_{46} and A_{56} ; all the other opinions A_{ij} serve as ‘transport’, i.e. merely providing weights multiplying the evidence in A_{46} and A_{56} . Therefore, there is no double counting of evidence.

In the case of Figure 3, i.e. with a loop, recursive use of (32) gives the direct EBSL equivalent of (16). We have

$$R_{13} = R_{12} \boxtimes A_{23} \quad (34)$$

$$\begin{aligned} R_{12} &= A_{12} \oplus R_{13} \boxtimes A_{32} \\ &= A_{12} \oplus (R_{12} \boxtimes A_{23}) \boxtimes A_{32}. \end{aligned} \quad (35)$$

Eq. (35) gives an expression for the unknown R_{12} that contains R_{12} . It can be solved in two ways. The first is to repeatedly substitute (35) into itself. We define a mapping $f(x) =$

$A_{12} \oplus (x \boxtimes A_{23}) \boxtimes A_{32}$. From (35) we see that R_{12} is a fixed point of f . The fixed point can be found approximately by setting $x = A_{12}$ and repeatedly applying f until the output does not change any more.

$$R_{12} \approx f^r(A_{12}) \quad \text{for large } r. \quad (36)$$

$$R_{12} = A_{12} \oplus ([A_{12} \oplus (\{A_{12} \oplus \dots\} \boxtimes A_{23}) \boxtimes A_{32}] \boxtimes A_{23}) \boxtimes A_{32}. \quad (37)$$

Convergence will be fast if A_{23} and A_{32} contain a lot of uncertainty. In Eq. (37) we observe that (i) The evidence contributing to R_{12} comes from the first A_{12} and from the final A_{32} . (ii) All the other occurrences of opinions A_{ij} are only used to compute the weights for the scalar multiplication.

The second method is to treat (35) as two independent equations in two unknowns (the two independent components of $R_{12} \in \Omega'$) and to solve them algebraically. This can be quite difficult if the function g is complicated, since g has to be applied twice,

$$R_{12} = A_{12} \oplus g(g(R_{12}) \cdot A_{23}) \cdot A_{32} \quad (38)$$

$$\Leftrightarrow \begin{cases} p(R_{12}) = p(A_{12}) + g(g(R_{12}) \cdot A_{23})p(A_{32}) \\ n(R_{12}) = n(A_{12}) + g(g(R_{12}) \cdot A_{23})n(A_{32}). \end{cases} \quad (39)$$

The solution is simple in the following special cases:

- If $g(A_{12}) = 0$ and $g(U) = 0$ then $R_{12} = A_{12}$.
- If $A_{23} \neq U$, $g(x) > 0$ for $x \neq U$, and $A_{32} \rightarrow B$ then $R_{12} \rightarrow B$. The direct A_{12} can be overwhelmed by the indirect A_{32} , even if user 1 has little trust in user 2. This demonstrates the danger of allowing opinions close to full belief.
- If $A_{23} \rightarrow B$ and $g(B) = 1$ then $R_{12} \rightarrow A_{12} \oplus A_{32}$.

6.2 Recursive solution in matrix form

The recursive equation (32) for obtaining the R_{ij} solutions can be rewritten in matrix notation. We choose g such that $g(U) = 0$. We are looking for the off-diagonal components of a matrix R or, equivalently, for a complete matrix R which has the uncertainty ‘ U ’ on its diagonal. Let X be an $n \times n$ matrix containing opinions; it is allowed to have a non-empty diagonal, e.g. $X_{ii} \neq U$. We define a function f as

$$f(X) \stackrel{\text{def}}{=} A \oplus \{\text{offdiag } X\} \boxtimes A, \quad (40)$$

where ‘offdiag X ’ is defined as X with its diagonal replaced by U entries, and an expression of the form $X \boxtimes A$ is a matrix defined as $(X \boxtimes A)_{ij} = \bigoplus_k X_{ik} \boxtimes A_{kj}$. Note that $f(X)$ is a matrix that can have a non-empty diagonal. Solving (32) is equivalent³ to the following procedure:

1. Find the fixed point X_* satisfying $f(X_*) = X_*$.
2. Take $R = \text{offdiag } X_*$.

³ An even more compact formulation is possible if one is willing to temporarily use the full belief B in the computations. Let ‘ $\mathbf{1}$ ’ be the diagonal matrix. Let $Z = B\mathbf{1} + R$. Solving (32) for R is equivalent to solving $Z = B\mathbf{1} \oplus (Z \boxtimes A)$ for Z .

One approach to determine the fixed point X_* is to pick a suitable starting matrix X_0 and then keep applying f until the output does not change any more, $X_* \approx f^N(X_0)$. Another approach is to treat the formula $f(X_*) = X_*$ as a set of algebraic equations, whose complexity is affected by the choice of the g function.

At this point, two important questions have to be answered: *i*) whether the recursive approach for (40) converges, i.e. if the fixed point exists, and *ii*), when it converges, whether the fixed point solution is unique. When there are no loops in the network, then trivially we have convergence and the fixed point is unique. Intuitively, the repeated applications of f after the trust network has been completely explored do not propagate additional evidence.

In the case of general networks the situation is more complicated. We can prove that there is *no divergence*. In every iteration of the mapping $X \mapsto f(X)$, the new value of X is of the form $X_{ij} = A_{ij} \oplus \bigoplus_k g((\text{offdiag } X)_{ik}) \cdot A_{kj}$. We see that the evidence in each A_{kj} gets multiplied by a scalar smaller than 1. Hence, no matter how many iterations are done, the amount of evidence about user j that is contained in X can never exceed the amount of evidence about j present in A . This puts a hard upper bound on the amount of evidence in $\text{offdiag } X$, which prevents the solution from ‘running off’ towards full Belief. Hence, the evidence underlying R_{ij} cannot be greater of the total amount of evidence underlying the opinions in A about user j .

It can be observed that, being flow-based, our fixed point equation for the matrix R has the same form as the fixed point equation for a Markov Chain. The main difference is that in an ordinary Markov chain there is a real-valued transition matrix whereas we have opinions $A_{ij} \in \Omega^*$, and in our case multiplication of reals is replaced by \boxtimes and addition by \oplus . In spite of these differences, we observe in our experiments that every type of behavior of Markov Chain flow also occurs for R . Indeed, experiments on real data show that we indeed have convergence (see Section 7). Moreover, for some fine-tuned instances of the A -matrix, which are exceedingly unlikely to occur naturally, oscillations can exist just like in Markov chains; after a number of iterations the powers of A jump back and forth between two states. Just as in flow-based reputation (1), adding the direct opinion matrix A in each iteration dampens the oscillations and causes convergence.

6.3 Recursive solutions using the \odot discounting rule

We investigate what happens when we replace the generic EBSL discounting operation \boxtimes by the special choice \odot as specified in Def. 13. First we consider the case of Figure 2. The generic Eq. (33) reduces to

$$\begin{aligned} R_{16} &= A_{12} \odot A_{23} \odot A_{34} \odot A_{46} \oplus \\ &\quad A_{12} \odot A_{23} \odot A_{35} \odot A_{56} \oplus \\ &\quad A_{12} \odot A_{23} \odot A_{34} \odot A_{45} \odot A_{56} \\ &= A_{12} \odot A_{23} \odot \{A_{34} \odot A_{46} \oplus (A_{35} \oplus A_{34} \odot A_{45}) \odot A_{56}\}. \end{aligned} \quad (41)$$

Notice that we do not have to put brackets around chains of \odot operations because they are associative (Lemma 9). Also notice that in (41) the common ‘factor’ $A_{12} \odot A_{23}$ has been pulled outside the brackets. We are allowed to add and remove brackets at will because of the associativity and full distributivity of \odot .

Next we consider Figure 3, the loop case. Eq. (39) reduces to

$$\begin{cases} p(R_{12}) = p(A_{12}) + \frac{1}{\bar{\theta}^2} p(R_{12}) p(A_{23}) p(A_{32}) \\ n(R_{12}) = n(A_{12}) + \frac{1}{\bar{\theta}^2} p(R_{12}) p(A_{23}) n(A_{32}) \end{cases} \quad (42)$$

which is easily solved,

$$p(R_{12}) = \frac{p(A_{12})}{1 - \theta^{-2}p(A_{23})p(A_{32})} \quad (43)$$

$$n(R_{12}) = n(A_{12}) + \frac{p(A_{12})p(A_{23})}{\theta^2 - p(A_{23})p(A_{32})}n(A_{32}). \quad (44)$$

We observe that the system parameter θ has to be chosen with great caution. If values $p(A_{ij})$ can get too close to θ then the fraction in (43) may explode and may result in $p(R_{12}) > \theta$, which is problematic. Let us define $p_{\max} = \max_{ij} p(A_{ij})$. Then it is necessary to set $\theta \geq p_{\max} \cdot (1 + \sqrt{5})/2$. (This bound is obtained by setting $p(A_{12}) = p(A_{23}) = p(A_{32}) = p_{\max}$ in (43) and demanding that $p(R_{12}) \leq \theta$.)

7 Evaluation

We have implemented our flow-based reputation model with uncertainty and performed a number of experiments to evaluate the practical applicability and “accuracy” of the model using both synthetic and real-life data. Note that, while it is possible to define some limiting situations (synthetic data) in which a certain result is expected, in general numerical experiments cannot ‘prove’ which approach is right because there is no ‘ground truth’ solution to the reputation problem that we could compare against. The only thing that can be verified by numerics is: (i) do the results make sense? (ii) is the method practical? Thus, we have used synthetic data to compare the accuracy of the opinions computed using different reputation models. On the other hand, we used real-life data to study the practical applicability.

Experiments for assessing the robustness of the reputation model against attacks like slander, self-promotion and Sybil attacks [8, 13, 35], have not been considered in this work and are left for future work, as our goal here is the definition of the mathematical foundation for the development of reputation systems. A study of the robustness against attacks requires to consider many other aspects that are orthogonal to this work.

In the remainder of this section, first we briefly present the implementation; then we report and discuss the results of the experiments.

7.1 Implementation

We have developed a tool in Python. It implements the procedure for computing the fixed point described in Section 6.2. All SL and EBSL computation rules presented in this paper have been implemented in a Python library.

The tool takes as input a log containing recorded user interactions. Based on the evidence contained in the log, the tool extracts the direct referral trust matrix A . The tool repeatedly iterates the recursive equation until it converges, that is, the difference between the new matrix $R^{(k+1)}$ and the previous one $R^{(k)}$ is less than a certain threshold. In particular the termination condition is set as follows:

$$\sum_{i,j} \delta(R_{ij}^{(k+1)}, R_{ij}^{(k)}) < 10^{-10} \quad (45)$$

where $\delta(x, y) \stackrel{\text{def}}{=} |x_b - y_b| + |x_d - y_d| + |x_u - y_u|$.

Source	Target	Evidence	Opinion	A_{xy}
1	2	(400, 300)	(0.570, 0.427, 0.003)	0.571
2	3	(10, 5)	(0.588, 0.294, 0.118)	0.667
3	4	(500, 0)	(0.996, 0.000, 0.004)	1.0
3	5	(500, 0)	(0.996, 0.000, 0.004)	1.0
4	5	(500, 0)	(0.996, 0.000, 0.004)	1.0
4	6	(500, 0)	(0.996, 0.000, 0.004)	1.0
5	6	(500, 0)	(0.996, 0.000, 0.004)	1.0
6	7	(5, 5)	(0.417, 0.417, 0.166)	0.5
7	P	(10, 90)	(0.098, 0.882, 0.020)	0.1

Table 1: Evidence, opinions, and aggregate ratings for case C1.

7.2 Synthetic Data

We have conducted a number of experiments using synthetic data to analyze and compare the different approaches for trust computation. The goal of these experiments is to analyze the behavior of the reputation models in a number of limiting situations for which it is known a priori how the result should behave.

Experiment Settings The experiments are based on the trust network in Figure 2. We considered six approaches: (i) the flow-based method without uncertainty in Eq. (1); (ii) the flow-based SL approach presented in Section 3; (iii) SL in which the specification of the trust network is transformed to canonical form by removing the edge from 4 to 5 (i.e., A_{45} is set to U in Eq. (14)); (iv) EBSL with $g(x) = x_b$; (v) EBSL with $g(x) = \sqrt{x_b}$; and (vi) EBSL using the operator \odot .

To make the results comparable, we specify the amount of positive and negative evidence (p, n) for each edge in the trust network and use such evidence to compute the opinions used for EBSL and SL as well as the aggregate ratings (A_{xy}) used for flow-based reputation (without uncertainty). The mapping between evidence and opinion is computed using (2). For the mapping between evidence and aggregated ratings, we use an approach similar to the one presented in [35]:

$$A_{xy} = \frac{1}{2} + \frac{p - n}{2e} \quad (46)$$

where $A_{xy} = 1$ means fully trusted, $A_{xy} = 0$ fully distrusted, and $A_{xy} = \frac{1}{2}$ neutral.

For the analysis we consider three cases. Table 1 presents the evidence along with the derived opinions and aggregated ratings for our first case (C1). In this case, we use $\theta = 1000$ for EBSL using operator \odot . In the second case (C2), we consider the same evidence except for the edge from 7 to P which is now (10, 900). The corresponding opinion and aggregate rating are $T_{7P} = (0.011, 0.987, 0.002)$ and $A_{7P} = 0.011$ respectively. In the last case (C3), we consider the evidence for all edges to be (10000, 0) and we set $\theta = 20000$. In this case, all opinions are equal to (0.9998, 0.0000, 0.0002) and aggregated ratings are equal to 1.

Results The results of the trust computation are presented in Table 2 in terms of opinions (trust value for flow-based method), and in Table 3 in terms of amount of evidence. Note that in Table 3 we have not included the amount of evidence for the flow-based approach of (1) as it is not possible to reconstruct it from trust values.

The results confirm our expectation about the impact of the trust network representation on the trust computation when SL is used. As expected, the uncertainty component of opinion F_{1P} computed using SL is larger when the trust network in Figure 2 is represented

	C1	C2	C3
Flow-based	0.401	0.392	0.501
Flow-SL	(0.024, 0.220, 0.756)	(0.003, 0.246, 0.751)	(0.9993, 0.0000, 0.0007)
SL (canonical form)	(0.014, 0.123, 0.863)	(0.002, 0.137, 0.861)	(0.9990, 0.0000, 0.0010)
EBSL ($g(x) = x_b$)	(0.095, 0.859, 0.046)	(0.011, 0.984, 0.005)	(0.9998, 0.0000, 0.0002)
EBSL ($g(x) = \sqrt{x_b}$)	(0.097, 0.873, 0.030)	(0.011, 0.986, 0.003)	(0.9998, 0.0000, 0.0002)
EBSL (\odot)	(0.000, 0.000, 1.000)	(0.000, 0.006, 0.994)	(0.9970, 0.0000, 0.0030)

Table 2: Comparison in terms of trust values r_{1P} and opinions F_{1P} .

	C1	C2	C3
Flow-SL	(0.065, 0.582)	(0.007, 0.655)	(2763.7, 0)
SL (canonical form)	(0.032, 0.284)	(0.004, 0.319)	(1999.16, 0)
EBSL ($g(x) = x_b$)	(4.166, 37.490)	(4.166, 374.901)	(9998, 0)
EBSL ($g(x) = \sqrt{x_b}$)	(6.455, 58.095)	(6.455, 580.946)	(9999, 0)
EBSL (\odot)	(0.000, 0.001)	(0.000, 0.011)	(781.25, 0)

Table 3: Comparison in terms of amount of evidence underlying opinion F_{1P} .

in canonical form. Indeed, in this case some trust information (i.e., evidence) is discarded (Recall that uncertainty depends on the amount of evidence; the larger the amount of evidence, the lower the uncertainty. See Section 5.2). In contrast, the representation of the trust network in Eq. (13) is affected by double counting, leading to more evidence and thus to a lower uncertainty component of F_{1P} .

The results show that the SL and EBSL approaches preserve the ratio between belief and disbelief components (Table 2) and consequently the ratio between positive and negative evidence (Table 3). This ratio is close to the one between the positive and negative evidence underlying the functional trust T_{7P} . If the amount of evidence increases (**C2**), one would expect that the amount of evidence underlying opinion F_{1P} increases proportionally to the increase of the amount of evidence underlying T_{7P} (Theorem 2). Accordingly, the amount of positive evidence underlying F_{1P} should be the same in **C1** and **C2**, and the amount of negative evidence underlying F_{1P} in **C2** should be ten times the amount of negative evidence in **C1**. We can observe in Table 3 that this is true for EBSL but not for SL. This is explained by the fact that $x \otimes y$ is not an x -dependent multiple of the evidence underlying y , as was shown in Eqs. (8) and (9).

Finally, in the last case (**C3**) we have considered a limiting case where every trust relation in the network of Figure 2 is characterized by a large amount of positive evidence. Here, one would expect that the opinion F_{1P} is close to $(1, 0, 0)$ and the trust value r_{1P} close to 1. From Table 2 we can observe that SL and all EBSL approaches meet this expectation. However, if we look closely at the evidence underlying such an opinion (Table 3), we can observe that when the SL discounting operator \otimes and EBSL operator \odot are used, a large amount of evidence is “lost” on the way. In contrast, we expect the amount of evidence underlying F_{1P} to be close to that of T_{7P} (Theorem 2). Table 3 shows that EBSL, both for $g(x) = x_b$ and $g(x) = \sqrt{x_b}$, preserves the amount of evidence when referral trust relations are close to full belief.

Moreover, Table 2 shows that the value of r_{1P} is close to neutral trust rather than to full trust. This can be explained by Eq. (1) and the impossibility to express uncertainty. On the one hand, at each iteration Eq. (1) computes a weighted average of aggregated ratings where weights are equal to the trust a user places in the users providing recommendations. On the other hand, the flow-based approach does not distinguish between neutral trust (equal

amount of positive and negative evidence) and full uncertainty (zero evidence). In particular, the lack of evidence between two users is represented in the matrix of aggregated ratings A as neutral trust (see Eq. 46). In sparse trust networks (i.e., networks with only a few edges) like the one in Figure 2, the neutral trust used to express uncertainty has a significant impact on the weighted average used to compute trust values.⁴ These results demonstrate that the ability to express uncertainty is fundamental to capture the actual trustworthiness of a target, which is one of the main motivations for this work.

7.3 Real-life Data

We performed a number of experiments using real-life data to assess the practical applicability of EBSL and flow-based reputation models built on top of EBSL. In particular, we study the impact of various discounting operators on the propagation of evidence and the convergence speed of the iterative procedure.

Experiment Settings For the experiments we used a dataset derived from a BitTorrent-based client called Tribler [29]. The dataset consists of information about 10,364 nodes and 44,796 interactions between the nodes. Each interaction describes the amount of transferred data in bytes from a node to another node. The amount of transferred data can be either negative, indicating an upload from the first node to the second node, or positive, indicating a download.

To provide an incentive for sharing information, some BitTorrent systems require users to have at least a certain ratio of uploaded vs. downloaded data. Along this line, we treat the amount of data uploaded by a user as positive evidence, and the downloaded amount as negative evidence. Intuitively, positive evidence indicates a user's inclination to share data and thus to contribute to the community.

It is worth noting that Tribler has a high population turnover and, thus, the dataset contains very few long living and active nodes alongside many loosely connected nodes of low activity [9]. This results in a direct referral trust matrix that is very sparse (i.e., most opinions are full uncertainty). In this sparse form it is inefficient to do large matrix multiplications. To this end, we have grouped the nodes into 200 clusters, each of which contains about 50 nodes. Intuitively, a cluster may be seen as the set of nodes under the control of a single user.

For the experiments with real data, we considered four reputation models: (i) the flow-based SL approach presented in Section 3; (ii) EBSL with $g(x) = x_b$; (iii) EBSL with $g(x) = \sqrt{x_b}$; and (iv) EBSL using the operator \odot . Note that, due to the large number of interconnected loops in the trust network, we did not consider SL in which the trust network is transformed into a canonical form.

In all four models we computed the final referral trust matrix R . The amount of evidence in the matrix A is visualized in Figure 4. Figure 4a presents the amount of positive evidence, and Figure 4b the total amount of evidence (sum of positive and negative). We can observe the presence of a few active users who had interactions with a lot of other users (visible as dark lines). A horizontal dark line in Figure 4a indicates a user who downloaded data from many other users. The vertical dark lines in Figure 4b represent negative evidence: many users uploading to the same few users. Note that Figure 4b is not symmetric, since an interaction never results in user feedback from both sides.

⁴ It worth noting that a trust value close to 1 can be obtained by Eq. 1 only if the trust network is a complete graph in which the aggregate rating associated with each edge is close to 1.

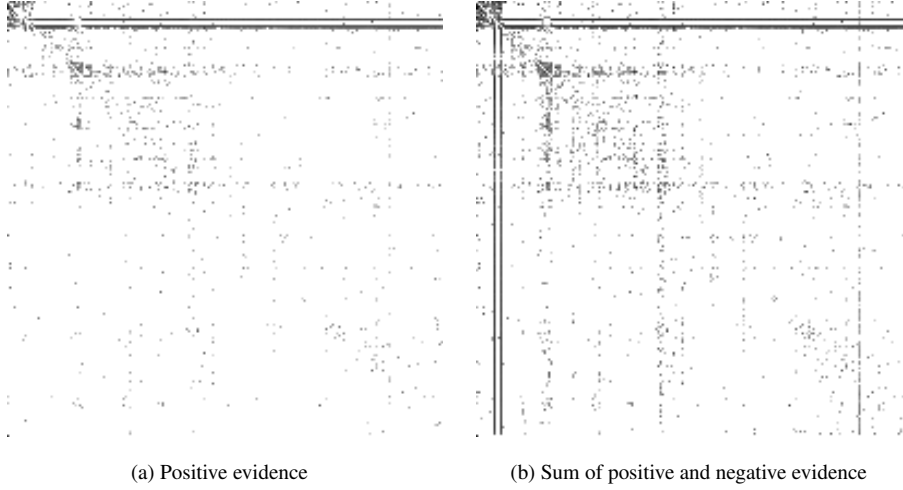


Fig. 4: The evidence in the direct opinion matrix A for the Tribler data. For each pair (i, j) the amount of evidence underlying the opinion of i about j is shown as a shade of gray, using a logarithmic gray scale. White corresponds to zero, black to $8.7 \cdot 10^6$, which is the maximum amount of evidence occurring in a single matrix entry in the experiments.

It is also interesting to note the clusters of strongly connected users who often interact with each other (for instance the top-left corner).

Results The amount of evidence in R is presented in Figure 5 (only positive evidence) and in Figure 6 (sum of positive and negative evidence). For most users, the amount of evidence in R has increased compared to the initial situation (Figures 4a and 4b). The plots are characterized by uniform vertical stripes, indicating that (most) users have approximately the same amount of (positive) evidence about a given user. The amount of evidence, however, remains close to 0 for those users who had very few interactions with other users (horizontal white lines in Figures 5 and 6). It is also worth noting that the diagonal of R is clearly recognizable as a white line. This is due to the fact we impose the diagonal to be full uncertainty, i.e. users cannot have an opinion about themselves, to reduce the effect of self-promoting.

The choice of the discounting operator, which defines how evidence is propagated, has a significant impact on the amount of evidence in R . Ideally, users should be able to use the available trust information to decide whether to engage an interaction with another user [42]. Therefore, a reputation system should allow users to gather as many recommendations (i.e., evidence) as possible from trusted users. However, the use of the \odot operator causes most of the evidence to be lost along the way. This can be clearly understood by observing that the initial situation in Figure 4a (Figure 4b respectively) and the final referral trust matrix R in Figure 5a (Figure 6a respectively) are almost the same. Figures 5b and 6b show that the \otimes operator propagates more evidence than \odot . We remind the reader that \otimes causes double-counting as well as discarding of evidence as shown in Examples 2 and 3. The \boxtimes operator with both $g(x) = x_b$ and $g(x) = \sqrt{x_b}$ results in the propagation of more evidence compared to the \odot and \otimes operators, as shown in Figures 5c and 5d (positive evidence) and in Figures 6c and 6d (total evidence). These findings confirm the results obtained with the synthetic data (Table 3). Therefore, we conclude that the \boxtimes operator is preferable to the other operators.

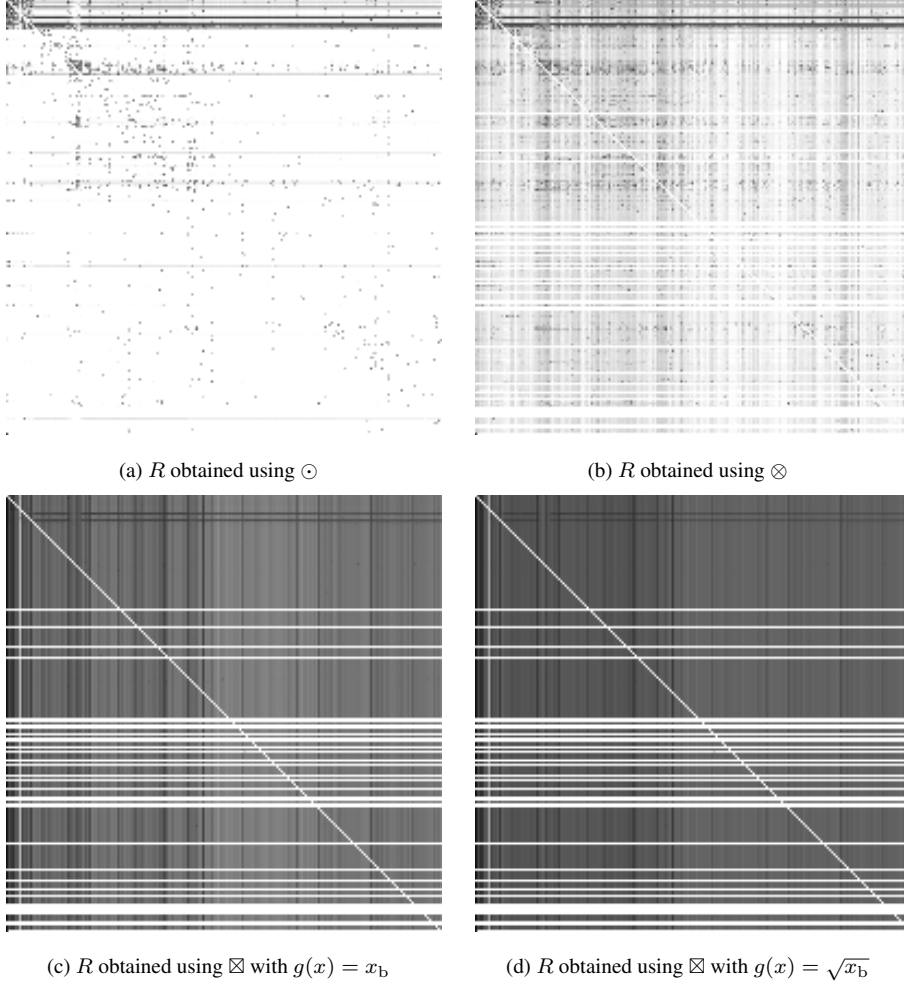


Fig. 5: Positive evidence in the final referral trust matrix R for the Tribler data. For each pair (i, j) the amount of positive evidence underlying the opinion of i about j is shown as a shade of gray, using a logarithmic gray scale. White corresponds to zero, black to $8.7 \cdot 10^6$, which is the maximum amount of evidence occurring in a single matrix entry in all the experiments.

Convergence We have analyzed the convergence of the iterative approach using the Tribler dataset. The experiments show that the reputation models built on top of EBSL converge. In particular, EBSL with $g(x) = x_b$ converges after 47 iterations, EBSL with $g(x) = \sqrt{x_b}$ converges after 24 iterations, and EBSL using the \odot operator after 9 iterations. One can observe that, in all the cases, convergence is fast. Accordingly, we believe that the proposed reputation model can handle real scenarios. In the experiments we also analyzed the convergence of the naïve approach that combines flow-based reputation and SL, as presented in Section 3. Here convergence is *not* reached in a reasonable amount of time: after 1000 iterations we still have $\sum_{i,j} \delta(R_{ij}^{(k+1)}, R_{ij}^{(k)}) \approx 10^{-8}$.

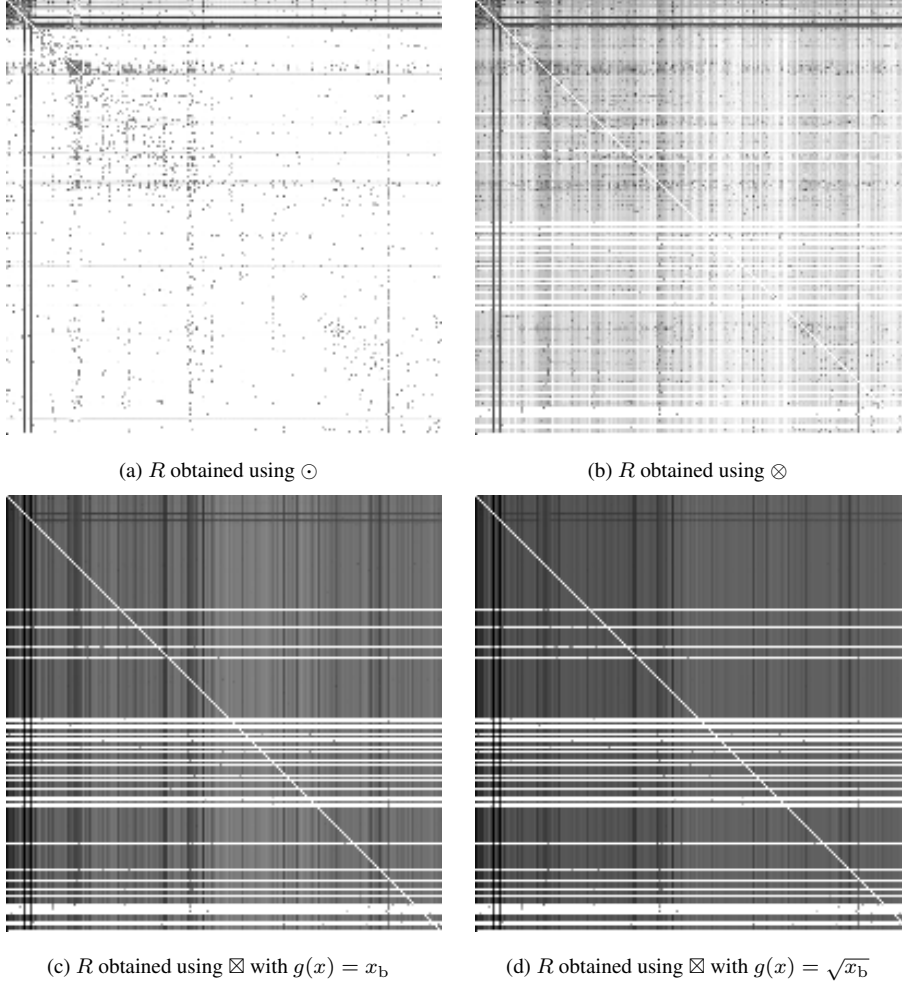


Fig. 6: The sum of positive and negative evidence in the final referral trust matrix R for the Tribler data. For each pair (i, j) the total amount of evidence underlying the opinion of i about j is shown as a shade of gray, using a logarithmic gray scale. White corresponds to zero, black to $8.7 \cdot 10^6$, which is the maximum amount of evidence occurring in a single matrix entry in the experiments.

To study the link between our approach and Markov chains, we performed additional experiments (not reported here) with a number of limiting situations, i.e. synthetic data unlikely to occur in real life. In particular, we studied the EBSL case where the powers of A show oscillations, i.e. $A^{k+2} = A^k$ with $A^{k+1} \neq A^k$. Here A^k stands for $((\dots \boxtimes A) \boxtimes A) \boxtimes A$. This can occur in Markov chains too. In flow-based reputation (1) the added term $(1 - \alpha)s_x$ dampens the oscillations and thus improves convergence. Similarly, in our EBSL experiments the added term A in each iteration (40) gives a convergent result in spite of the oscillatory nature of A . This strengthens our conviction that EBSL correctly captures the idea of reputation flow.

8 Related Work

The notion of uncertainty is becoming an important concept in reputation systems and, more in general, in data fusion [5, 24]. Uncertainty has been proposed as a quantitative measure of the accuracy of predicted beliefs and it is used to represent the level of confidence in the fusion result. Several approaches have extended reputation systems with the notion of uncertainty [15, 30, 32, 39]. For instance, Reis et al. [32] associate a parameter with opinions to indicate the degree of certainty to which the average rating is assumed to be representative for the future. Teacy et al. [39] account for uncertainty by assessing the reputation of information sources based on the perceived accuracy of past opinions. Differently from the previous approaches, Subjective Logic [15] considers uncertainty as a dimension orthogonal to belief and disbelief, which is based on the amount of available evidence.

One of the main challenges in reputation systems is how to aggregate opinions, especially in the presence of uncertainty. SL provides two main operators for aggregating opinions: consensus and discounting (see Section 2.2). Many studies have analyzed strategies for combining conflicting beliefs [16, 19, 37] and have proposed new combining strategies and operators [7, 44]. In Section 5.2, we re-confirm that the standard consensus operator used in SL is well-founded on the theory of evidence.

In contrast, less effort has been devoted to studying the discounting operator. Bhuiyan and Jøsang [4] propose two alternative discounting operators: an operator based on opposite belief favouring, for which the combination of two disbeliefs results in belief, and a base rate sensitive transitivity operator in which the trust in the recommender is a function of the base rate. Similarly to the traditional discounting operator of SL, these operators are founded on probability theory. As shown in Section 3, employing operators founded on different theories has the disadvantage that these operators may not “cooperate”. In the case of SL, this lack of cooperation results in the inability to apply SL to arbitrary trust networks. In particular, trust networks have to be expressed in a canonical form in which edges are not repeated. A possible strategy to reduce a trust network to a canonical form is to remove the weakest edges (i.e., the least certain paths) until the network can be expressed in canonical form [20]. This, however, has the disadvantage that some (possibly even much) trust information is discarded. An alternative canonicalization method called edge splitting was presented in [18]. The basic idea of this method is to split a dependent edge into a number of different edges equal to the number of different instances of the edge in the network expression. Nonetheless, the method requires that the trust network is acyclic; if a loop occurs in the trust network, some edges have to be removed in order to eliminate the loop, thus discarding trust information. In contrast, we have constructed a discounting operator founded on the theory of evidence. This operator together with the consensus operator allows the computation of reputation for arbitrary trust networks, which can include loops, without the need to discard any information.

Cerutti et al. [7] define three requirements for discounting based on the intuitive understanding of few scenarios: Let A be x ’s opinion about y ’s trustworthiness, C the level of certainty that y has about a proposition P , and $F = A \circ C$ the (indirect) opinion that x has about P . (i) If C is pure belief, then $F = A$; (ii) If C is complete uncertainty, then $F = C$; (iii) The belief part of F is always less than or equal to the belief part of A . Based on these requirements, they propose a family of graphical discounting operators which, given two opinions, project one opinion into the admissible space of opinions given by the other opinion. These operators are founded on geometric properties of the opinion space. This makes it difficult to determine whether the resulting theory is consistent with the theory of evidence or probability theory. Our discounting operator satisfies requirement (ii) above,

	\otimes	\boxtimes	\odot
Associativity	yes	no	yes
Left-distribution	no	no	yes
Right-distribution	no	yes	yes
Recursive solutions	no	yes	yes

Table 4: Comparison of the operators \otimes , \boxtimes , and \odot .

but not requirements (i) and (iii); indeed, for $g(A) > 0$ it holds that $A \boxtimes B = B$ (where B represents full belief). It is worth noting that the requirements proposed in [7] are not well founded in the theory of evidence: B means that there is an infinite amount of positive evidence; discounting an infinite amount of evidence still gives an infinite amount of evidence. In Theorem 1, we provided a number of desirable properties founded on the theory of evidence. In particular, if $p + n \rightarrow \infty$ then $u \rightarrow 0$. Accordingly, if $C = B$ the uncertainty component of F should be equal to 0, regardless of the precise (nonzero) value of the uncertainty component of A .

To our knowledge, our proposal is the first work that integrates uncertainty into flow-based reputation.

9 Conclusion

In this paper, we have presented a flow-based reputation model with uncertainty that allows the construction of an automated reputation assessment procedure for arbitrary trust networks. We illustrated and discussed the limitations of a naïve approach to combine flow-based reputation and SL. An analysis of SL shows that the problem is rooted in the lack of “cooperation” between the SL consensus and discounting rules due to the different nature of these two operators. In order to solve this problem, we have revised SL by introducing a scalar multiplication operator and a new discounting rule based on the flow of evidence. We refer to the new opinion algebra as Evidence-Based Subjective Logic (EBSL).

A generic definition of discounting (the operator \boxtimes) lacks the associative property satisfied by the SL operator \otimes . This, however, is not problematic since the flow of evidence has a well defined direction. Furthermore, the operator \boxtimes has right-distributivity, a property that one would intuitively expect of opinion discounting. One can choose a specific discounting function $g(x)$ proportional to the amount of positive evidence in x . The resulting discounting operator is denoted as \odot . As shown in Table 4, this operator is completely linear (associative as well as left and right distributive). However, it has potentially undesirable behavior since it ignores negative evidence, and requires a carefully chosen system parameter related to the maximum amount of positive evidence in the system.

The adoption of the discounting operator \boxtimes results in a system that is centered entirely on the handling of evidence. We have showed that this new EBSL algebra makes it possible to define an iterative algorithm to compute reputation for arbitrary trust networks. Thus, EBSL poses the basis for the development of novel reputation systems. In particular, our opinion algebra guarantees that trust information does not have to be discarded, thus preserving the quality of the aggregated evidence. Moreover, making the uncertainty of the computed reputation explicit helps users in deciding how much to rely on it based on their risk attitude. In our opinion, this will facilitate the adoption and acceptance of reputation systems since users are more aware of the risks of engaging a transaction and, thus, can make more informed decisions.

The work presented in the paper poses the basis for several directions of future work. We have shown how EBSL can be used to build a flow-based reputation model with uncertainty. However, several reputation models have been proposed in the literature to compute reputation over a trust network. An interesting direction is to study the applicability of EBSL as a mathematical foundation for these models. This will also make it possible to study the impact of uncertainty on the robustness of reputation systems against attacks like self-promotion, slander and Sybil attacks.

Acknowledgements This work has been partially funded by the Dutch national program COMMIT under the THeCS project, the ITEA2 project FedSS, the ARTEMIS project ACCUS, and the EDA project IN4STARS 2.0.

References

1. Abdul-Rahman, A., Hailes, S.: A Distributed Trust Model. In: Proceedings of the 1997 Workshop on New Security Paradigms, pp. 48–60. ACM (1997)
2. Asnar, Y., Zannone, N.: Perceived risk assessment. In: Proceedings of the 4th ACM Workshop on Quality of Protection, pp. 59–64. ACM (2008)
3. Bharadwaj, K., Al-Shamri, M.: Fuzzy computational models for trust and reputation systems. *Electron. Commer. Res. Appl.* **8**(1), 37–47 (2009)
4. Bhuiyan, T., Jøsang, A.: Analysing Trust Transitivity and The Effects of Unknown Dependence. *International Journal of Engineering Business Management* **2**(1), 23–28 (2010)
5. Bleiholder, J., Naumann, F.: Data Fusion. *ACM Comput. Surv.* **41**(1), 1:1–1:41 (2009)
6. Brin, S., Page, L.: The anatomy of a large-scale hypertextual Web search engine. *Computer networks and ISDN systems* **30**(1-7), 107–117 (1998)
7. Cerutti, F., Toniolo, A., Oren, N., Norman, T.J.: Subjective Logic Operators in Trust Assessment: an Empirical Study. *arxiv:1312.4828* (2013)
8. Douceur, J.R.: The Sybil Attack. In: Proc. of International Workshop on Peer-To-Peer Systems, LNCS 2429, pp. 251–260. Springer (2002)
9. Gkorou, D., Vinkó, T., Pouwelse, J.A., Epema, D.H.J.: Leveraging node properties in random walks for robust reputations in decentralized networks. In: Proceedings of International Conference on Peer-to-Peer Computing, pp. 1–10. IEEE (2013)
10. Govindan, K., Mohapatra, P.: Trust Computations and Trust Dynamics in Mobile Adhoc Networks: A Survey. *IEEE Communications Surveys Tutorials* **14**(2), 279–298 (2012)
11. Grandison, T., Sloman, M.: A Survey of Trust in Internet Applications. *Commun. Surveys Tuts.* **3**(4), 2–16 (2000)
12. Heuser, H.: *Funktionalanalysis: Theorie und Anwendung*, 4th edn. B.G. Teubner Verlag (2006)
13. Hoffman, K., Zage, D., Nita-Rotaru, C.: A Survey of Attack and Defense Techniques for Reputation Systems. *ACM Comput. Surv.* **42**, 1:1–1:31 (2009)
14. Huynh, T.D., Jennings, N.R., Shadbolt, N.R.: An integrated trust and reputation model for open multi-agent systems. *AAMAS* **13**(2), 119–154 (2006)
15. Jøsang, A.: A logic for uncertain probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **9**(3), 279–311 (2001)
16. Jøsang, A.: The consensus operator for combining beliefs. *Artificial Intelligence* **141**(12), 157–170 (2002)
17. Jøsang, A., Azderska, T., Marsh, S.: Trust Transitivity and Conditional Belief Reasoning. In: Trust Management VI, *IFIP Advances in Information and Communication Technology*, vol. 374, pp. 68–83. Springer (2012)
18. Jøsang, A., Bhuiyan, T.: Optimal Trust Network Analysis with Subjective Logic. In: Proceedings of 2nd International Conference on Emerging Security Information, Systems and Technologies, pp. 179–184. IEEE (2008)
19. Jøsang, A., Daniel, M., Vannoorenberghe, P.: Strategies for Combining Conflicting Dogmatic Beliefs. In: Proceedings of the 6th International Conference on Information Fusion (2003)
20. Jøsang, A., Gray, E., Kinatader, M.: Simplification and Analysis of Transitive Trust Networks. *Web Intelligence and Agent Systems* **4**, 139–161 (2006)
21. Jøsang, A., Haller, J.: Dirichlet reputation systems. In: Proc. of Int. Conf. on Availability, Reliability and Security, pp. 112–119. IEEE (2007)

22. Jøsang, A., Hayward, R., Pope, S.: Trust Network Analysis with Subjective Logic. In: Proceedings of the 29th Australasian Computer Science Conference, pp. 85–94. Australian Computer Society, Inc. (2006)
23. Kamvar, S., Schlosser, M., Garcia-Molina, H.: The EigenTrust algorithm for reputation management in P2P networks. In: Proceedings of the 12th International Conference on World Wide Web, pp. 640–651. ACM (2003)
24. Keenan, T., Carbone, M., Reichstein, M., Richardson, A.: The modeldata fusion pitfall: assuming certainty in an uncertain world. *Oecologia* **167**(3), 587–597 (2011)
25. Lempel, R., Moran, S.: The stochastic approach for link-structure analysis (SALSA) and the TKC effect. *Computer Networks* **33**(1), 387–401 (2000)
26. Lin, C., Varadharajan, V.: MobileTrust: a trust enhanced security architecture for mobile agent systems. *International Journal of Information Security* **9**(3), 153–178 (2010)
27. Liu, L., Munro, M.: Systematic analysis of centralized online reputation systems. *Decision Support Systems* **52**(2), 438–449 (2012)
28. Muller, T., Schweitzer, P.: On Beta Models with Trust Chains. In: Trust Management VII, IFIP AICT 401, pp. 49–65. Springer (2013)
29. Pouwelse, J.A., Garbacki, P., Wang, J., Bakker, A., Yang, J., Iosup, A., Epema, D.H.J., Reinders, M., van Steen, M.R., Sips, H.J.: TRIBLER: A Social-based Peer-to-peer System. *Concurrency and Computation: Practice and Experience* **20**(2), 127–138 (2008)
30. Regan, K., Cohen, R., Poupart, P.: The Advisor-POMDP: A Principled Approach to Trust through Reputation in Electronic Markets. In: Proceedings of 3rd Annual Conference on Privacy, Security and Trust (2005)
31. Richters, O., Peixoto, T.P.: Trust Transitivity in Social Networks. *PLoS ONE* **6**(4) (2011)
32. Ries, S., Habib, S.M., Mühlhäuser, M., Varadharajan, V.: CertainLogic: A Logic for Modeling Trust and Uncertainty. In: Trust and Trustworthy Computing, LNCS 6740, pp. 254–261. Springer (2011)
33. Sentz, K., Ferson, S.: Combination of evidence in Dempster-Shafer theory. Tech. rep., Sandia National Laboratories, Albuquerque, New Mexico (2002)
34. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press (1976)
35. Simone, A., Škorić, B., Zannone, N.: Flow-based reputation: more than just ranking. *International Journal of Information Technology & Decision Making* **11**(3), 551–578 (2012)
36. Smeltzer, L.R.: The Meaning and Origin of Trust in Buyer-Supplier Relationships. *Journal of Supply Chain Management* **33**(1), 40–48 (1997)
37. Smets, P.: Analyzing the combination of conflicting belief functions. *Information Fusion* **8**(4), 387–412 (2007)
38. Song, S., Hwang, K., Zhou, R., Kwok, Y.: Trusted P2P transactions with fuzzy reputation aggregation. *Internet Computing* **9**(6), 24–34 (2005)
39. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: TRAVOS: Trust and Reputation in the Context of Inaccurate Information Sources. *Autonomous Agents and Multi-Agent Systems* **12**(2), 183–198 (2006)
40. Trivellato, D., Zannone, N., Etalle, S.: GEM: A distributed goal evaluation algorithm for trust management. *TPLP* **14**(3), 293–337 (2014)
41. Vavilis, S., Petković, M., Zannone, N.: Impact of ICT on Home Healthcare. In: ICT Critical Infrastructures and Society, IFIP AICT 386, pp. 111–122. Springer (2012)
42. Vavilis, S., Petković, M., Zannone, N.: A reference model for reputation systems. *Decision Support Systems* **61**, 147–154 (2014)
43. Xiong, L., Liu, L.: PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities. *IEEE Trans. on Knowl. and Data Eng.* **16**(7), 843–857 (2004)
44. Zhou, H., Shi, W., Liang, Z., Liang, B.: Using new fusion operations to improve trust expressiveness of subjective logic. *Wuhan University Journal of Natural Sciences* **16**(5), 376–382 (2011)
45. Zomlot, L., Sundaramurthy, S.C., Luo, K., Ou, X., Rajagopalan, S.R.: Prioritizing Intrusion Analysis Using Dempster-Shafer Theory. In: Proceedings of 4th ACM Workshop on Security and Artificial Intelligence, pp. 59–70. ACM (2011)

A Associativity and Left-Distributivity of \boxtimes Imply Linearity of g

We present a proof of Lemma 10. The proof technique requires that opinions and the operations \oplus and \cdot form a vector space with an inner product defined on it. There is no such thing as an opinion ‘ $-x$ ’ (inverse of x with respect to \oplus) in Ω or Ω^I , so we have to broaden our scope. We define an extended opinion space

$$\Omega^* \stackrel{\text{def}}{=} \{(b, d, u) \in [-1, 1]^2 \times [0, 1] \mid |b| + |d| + u = 1\}$$

that allows for negative belief and/or disbelief components. We have to slightly modify the relation between evidence and opinions, so that negative *amounts* of evidence can be represented,

$$(b, d, u) = \frac{(p, n, c)}{|p| + |n| + c} \quad ; \quad (p, n) = c \frac{(b, d)}{u}$$

with $c > 0$. Here p and n can be negative. This relation automatically leads to a slightly modified definition of evidence addition (\oplus) and scalar multiplication,

$$\begin{aligned} x \oplus y &\stackrel{\text{def}}{=} \frac{(x_b y_u + y_b x_u, x_d y_u + y_d x_u, x_u y_u)}{|x_b y_u + y_b x_u| + |x_d y_u + y_d x_u| + x_u y_u} \\ \alpha \cdot x &\stackrel{\text{def}}{=} \frac{(\alpha x_b, \alpha x_d, x_u)}{|\alpha x_b| + |\alpha x_d| + x_u}. \end{aligned}$$

For $x, y \in \Omega$ and $\alpha \geq 0$ all this reduces to the algebra of Sections 2.2 and 5.3; For $x, y \in \Omega^*$ and $\alpha \in \mathbb{R}$ all the nice linear properties still hold.

The space Ω^* with the \oplus and \cdot operations is a vector space. (The underlying space of (p, n) evidence pairs has also been turned into a vector space by allowing negative amounts of evidence.) We introduce an inner product on this vector space as follows,

$$\langle x, y \rangle \stackrel{\text{def}}{=} p(x)p(y) + n(x)n(y).$$

It is easily verified from the definitions that this expression satisfies all the requirements for being an inner product, namely $\langle x, y \rangle = \langle y, x \rangle$; $\langle x, y \oplus z \rangle = \langle x, y \rangle + \langle x, z \rangle$; $\langle x, x \rangle \geq 0$ and $\langle \alpha \cdot x, y \rangle = \alpha \langle x, y \rangle$.

With all this structure in place we can now invoke the Riesz-Fréchet theorem [12],

If Ω^* is a real Hilbert space and g a linear functional, then there exist a unique vector $v \in \Omega^*$ such that $g(x) = \langle v, x \rangle$ for all $x \in \Omega^*$.

Here ‘linear functional’ means that the linear property $g(x \oplus y) = g(x) + g(y)$ holds. Hence, the only way to achieve this property is to set $g(x) = \langle v, x \rangle$, i.e. a linear combination of $p(x)$ and $n(x)$.